The Impact of Social Ignorance on Weighted Congestion Games*

Dimitris Fotakis[1], Vasilis Gkatzelis[2], Alexis C. Kaporis[3,4], and Paul G. Spirakis[4]

[1] School of Electrical and Computer Engineering
National Technical University of Athens, 15780 Athens, Greece
[2] Computer Science Department, Courant Institute
New York University, 251 Mercer Street, NY 10012, USA
[3] Department of Information and Communication Systems Engineering
University of the Aegean, 83200 Samos, Greece
[4] Research Academic Computer Technology Institute
N. Kazantzaki Str., University Campus, 26500 Patras, Greece

Emails: fotakis@cs.ntua.gr, gkatz@cims.nyu.edu,
kaporis@ceid.upatras.gr, spirakis@cti.gr

Abstract. We consider weighted linear congestion games, and investigate how social ignorance, namely lack of information about the presence of some players, affects the inefficiency of pure Nash equilibria (PNE) and the convergence rate of the ε-Nash dynamics. To this end, we adopt the model of graphical linear congestion games with weighted players, where the individual cost and the strategy selection of each player only depends on his neighboring players in the social graph. We show that such games admit a potential function, and thus a PNE. Our main result is that the impact of social ignorance on the Price of Anarchy (PoA) and the Price of Stability (PoS) is naturally quantified by the *independence number* $\alpha(G)$ of the social graph G. In particular, we show that the PoA grows roughly as $\alpha(G)(\alpha(G)+2)$, which is essentially tight as long as $\alpha(G)$ does not exceed half the number of players, and that the PoS lies between $\alpha(G)$ and $2\alpha(G)$. Moreover, we show that the ε-Nash dynamics reaches an $\alpha(G)(\alpha(G)+2)$-approximate configuration in time that is polynomial and does not directly depend on the social graph. For unweighted graphical linear games with symmetric strategies, we show that the ε-Nash dynamics converges to an ε-approximate PNE in time that is polynomial and exceeds the corresponding time for symmetric linear games by a factor at most as large as the number of players.

* This work was partially supported by the ICT Programme of the EU, under contract no. ICT-2008-215270 (FRONTS), by NSF grant CCF0830516, and by Andreas Mentzelopoulos Scholarships for the University of Patras.

1 Introduction

Congestion games provide a natural model for non-cooperative resource allocation in large-scale systems and have been the subject of intensive research in algorithmic game theory. In a (weighted) *congestion game*, a finite set of non-cooperative players, each controlling an unsplittable (weighted) demand, compete over a finite set of resources. All players using a resource experience a delay (or cost) given by a non-negative and non-decreasing function of the resource's total demand (or congestion). Among a given set of resource subsets (or strategies), each player selects one selfishly trying to minimize his *individual cost*, that is the sum of the delays on the resources in the chosen strategy. A natural solution concept is that of a *pure Nash equilibrium* (PNE), a configuration where no player can decrease his individual cost by unilaterally changing his strategy. Rosenthal [20] proved that the PNE of (unweighted) congestion games correspond to the local optima of a natural potential function, and thus every congestion game admits a PNE. A similar result was shown more recently for weighted congestion games with linear resource delays [15].

Motivation and Previous Work. The prevailing questions in recent work on congestion games have to do with quantifying the inefficiency due to the players' non-cooperative and selfish behaviour (see e.g. [19, 3, 5, 12, 11, 2, 9, 13]), and bounding the convergence time to (approximate) PNE if the players select their strategies in a selfish and decentralized fashion (see e.g. [14, 1, 10, 21, 6]).

Inefficiency of Pure Nash Equilibria. It is well known that a PNE may not optimize the system performance, usually measured by the *total cost* incurred by all players. The main tools for quantifying and understanding the performance degradation due to the players' non-cooperative and selfish behaviour have been the *Price of Anarchy* (PoA), introduced by Koutsoupias and Papadimitriou [19], and the *Price of Stability* (PoS), introduced by Anshelevich *et al.* [3]. The (pure) PoA (resp. PoS) is the *worst-case* (resp. *best-case*) ratio of the total cost of a PNE to the optimal total cost.

Many recent contributions have provided tight upper and lower bounds on the PoA and the PoS for several interesting classes of congestion games, mostly congestion games with linear and polynomial delays. Awerbuch *et al.* [5] and Christodoulou and Koutsoupias [12] proved that the PoA of congestion games is $5/2$ for linear delays and $d^{\Theta(d)}$ for polynomial delays of degree d. Subsequently, Aland *et al.* [2] obtained exact bounds on the PoA for congestion games with polynomial delays. For weighted congestion games with linear delays, Awerbuch *et al.* [5] proved that the PoA is $(3 + \sqrt{5})/2$. Christodoulou and Koutsoupias and Caragiannis *et al.* [11, 9] proved that the PoS for congestion games with linear delays is $1 + \sqrt{3}/3$. Recently, Christodoulou *et al.* [13] obtained tight bounds on the PoA and the PoS of approximate PNE for congestion games with linear delays.

Convergence Time to Pure Nash Equilibria. The existence of a potential function implies that a PNE is reached in a natural way when players iteratively select strategies that improve on their individual cost, given the strategies of other players. Nevertheless, this may take an exponential number of steps, since computing a PNE is PLS-complete even for symmetric congestion games and for asymmetric network games with linear delays [14, 1]. In fact, the proofs of [14, 1] establish the existence of instances where any sequence of players' improvement moves is exponentially long.

A natural approach to circumvent the strong negative results of [14, 1] is to resort to *approximate* PNE, where no player can *significantly* improve his individual cost by unilaterally changing his strategy. Chien and Sinclair [10] considered symmetric congestion games with a weak restriction on the delay functions, and proved that several natural families of sequences of significant improvement moves converge to an approximate PNE in polynomial time. On the other hand, Skopalik and Vöcking [21] proved that computing an approximate PNE for asymmetric congestion games is PLS-complete, and that even with the restriction of [10] on the delay functions, there are instances where any sequence of significant improvement moves leading to an approximate PNE is exponentially long. Nevertheless, Awerbuch *et al.* [6] showed that for unweighted congestion games with polynomial delays and for weighted games with linear delays, many natural families of sequences of significant improvement moves reach an approximately optimal configuration in polynomial time, where the approximation ratio is arbitrarily close to the PoA of the game.

Social Ignorance in Congestion Games. Most of the recent work on congestion games (including all the references above) focuses on the full information setting, where each player knows the precise weights and the actual strategies of all other players, and his strategy selection takes all this information into account. In many typical applications of congestion games however, the players arguably have incomplete information not only about the precise weights or the actual strategies, but also about the mere existence of (some of) the players with whom they compete for resources[5] (see e.g. [16, 18, 17, 7, 8] for similar considerations). In fact, in many applications of congestion games, it is both natural and convenient to assume that there is a *social context* associated with the game, and that the social context essentially determines the information available to the players. In particular, one may assume that each player has complete information about the set of players in his *social neighborhood*, and limited (if any) information about the other players.

The motivation of this work is to investigate how such social-context-related information considerations affect the inefficiency of PNE and the convergence rate to approximate PNE. To come up with a manageable setting that allows for some concrete answers, we make the simplifying assumption that each player has complete information about the players in his social neighborhood, and no information whatsoever about the remaining players. Therefore, since each player is not aware of the players outside his social neighborhood, his individual cost and his strategy selection are not affected by them. In fact, this is the model of *graphical congestion games*, introduced by Bilò, Fanelli, Flammini, and Moscardelli [7], and motivated by very similar considerations. The new ingredient in the definition of graphical congestion games is the *social graph*, which represents the players' social context. The social graph is defined on the set of players and contains an edge between each pair of players that know each other. The basic idea (and assumption) behind graphical congestion games is that the individual cost (aka *presumed cost*) of each player only depends on the players in his social neighborhood, and thus his strategy selection is only affected by them.

Bilò *et al.* [7] considered unweighted graphical congestion games, and proved that such games with linear delays and undirected social graphs admit a potential function, and thus a PNE. In addition, they proved that graphical games with directed acyclic social graphs have a PNE reached by a particular sequence of best response moves. For unweighted linear graphical games, Bilò *et al.* proved that the PoS is at most n, and the PoA is at most $n(\deg_{\max} + 1)$, where n is the number of players and $\deg_{\max}$ is the maximum degree of the social graph, and presented certain families of instances for which these bounds are tight. To the best of our understanding, the fact that these bounds are tight for some instances illustrates that expressing the PoA and the PoS as functions of n and $\deg_{\max}$ only does not provide an accurate picture of the impact of the social ignorance (see also the discussion in [7, Section 1.2]). In particular, the bound on the PoA conveys the message that the more the players know (or learn) about other players, the worse the PoA becomes, and fails to capture that as the social graph tends to the complete graph, the PoA should become a small constant that tends to $5/2$.

Contribution. Adopting graphical linear congestion games as a model, we investigate whether there is a natural parameter of the social graph that completely characterizes the impact of social ignorance

[5] In many applications, information considerations have to do not only with what the players actually know or are able to learn about the game, but also with how much information the players are able or willing to handle in their strategy selection process.

on the inefficiency of PNE and on the convergence rate of the ε-Nash dynamics. We restrict our attention to graphical linear games with *undirected* social graphs. We consider *weighted players*, so as to investigate the potential additional impact of different weights (i.e. could the PoA and the PoS become worse, and if yes, by how much, when many "small" players ignore a few "large" ones compared against the same social situation with all players of the same weight?). With a single exception, the PoA (resp. PoS) is defined with respect to the *actual total cost* of the worst (resp. best) PNE, while equilibria are defined with respect to the *presumed cost* (i.e. in a PNE, no player can improve his presumed cost, which is an underestimation of his actual cost based on limited social knowledge).

We prove that the impact of social ignorance on the PoA and the PoS is naturally quantified by the *independence number* $\alpha(G)$ of the social graph G, i.e. by the cardinality of the largest set of players that do not know each other. In particular, we show that the PoA grows roughly as $\alpha(G)(\alpha(G)+2)$, which is essentially tight as long as $\alpha(G) \leq n/2$, and that the PoS lies between $\alpha(G)$ and $2\alpha(G)$.

From a technical point of view, we first show that graphical linear games with weighted players admit a potential function, and thus a PNE (cf. Theorem 1). We remark that our potential function nicely generalizes the potential function of [15, Theorem 3.2], where the social graph is complete, and the potential function of [7, Theorem 1], where the players are unweighted.

To bound the PoA and the PoS from above, we show that the total actual cost in any configuration is an $\alpha(G)$-approximation of the total presumed cost of the players in the same configuration (cf. Lemma 1). Then, we prove that the PoA of any graphical linear congestion game with weighted players is at most $\alpha(G)(\alpha(G)+2+\sqrt{\alpha^2(G)+4\alpha(G)})/2$, which varies from $(3+\sqrt{5})/2$, when the social graph is complete[6], to roughly $\alpha(G)(\alpha(G)+2)$, when $\alpha(G)$ is large (cf. Theorem 2). Furthermore, we show that the upper bound is essentially tight, even for unweighted players, in the most interesting case where $\alpha(G)$ is no more than half the number of players (cf. Theorem 3). We also prove that the PoS is at most $\frac{2n\alpha(G)}{n+\alpha(G)}$ (cf. Theorem 6) and at least $\alpha(G)-\varepsilon$, for any $\varepsilon > 0$ (cf. Theorem 7). It is rather surprising that the upper bounds on the PoA and the PoS only depend on the *cardinality* of the largest set of players that do not know each other, not on their weights. In addition, the fact that all our lower bounds are proven for unweighted players implies that as long as the worst-case PoA and PoS are concerned, considering players with different weights cannot make things worse.

Moreover, we prove that the upper bound of $\alpha(G)(\alpha(G)+2+\sqrt{\alpha^2(G)+4\alpha(G)})/2$ holds even if the PoA is calculated with respect to the total presumed cost of the players (cf. Theorem 4), and that this bound is essentially tight as long as $\alpha(G) \leq \sqrt{n/2}$ (cf. Theorem 5).

As for the convergence time to approximately optimal configurations, we show that it does not directly depend on the structure of the social graph, only the approximation ratio does. In particular, using the techniques of Awerbuch *et al.* [6], we show that the *largest improvement* ε-Nash dynamics reaches an approximately optimal configuration in a polynomial number of steps (cf. Theorem 8). The approximation ratio is arbitrarily close to the PoA, so it is roughly $\alpha(G)(\alpha(G)+2)$, while the convergence time is linear in n and in the logarithm of the initial potential value, with the only dependence on the structure of the social graph hidden in the latter term. For graphical linear games with unweighted players and symmetric strategies, we use the techniques of Chien and Sinclair [10], and show that the largest improvement ε-Nash dynamics converges to an ε-PNE in a polynomial number of steps (cf. Theorem 9). Compared to the bound implied by [10, Theorem 3.1] for symmetric linear

[6] By a different analysis, we can show that the PoA for unweighted graphical linear games is at most $\frac{3\alpha(G)+7}{3\alpha(G)+1}\alpha^2(G)$. We omit the details, since $\frac{3\alpha(G)+7}{3\alpha(G)+1}\alpha^2(G)$ is better than $\alpha(G)(\alpha(G)+2+\sqrt{\alpha^2(G)+4\alpha(G)})/2$ only for $\alpha(G)=1$, in which case the PoA is $\frac{5}{2}$ for unweighted players and $\frac{3+\sqrt{5}}{2}$ for weighted players [5, 12], and for $\alpha(G)=2$, in which case the PoA is at most $\frac{52}{7} \approx 7.429$ for unweighted players and at most $4+\sqrt{12} \approx 7.464$ for weighted players.

games, the convergence time increases by a factor up to n due to the social ignorance. Both results can be extended to the unrestricted ε-Nash dynamics, which proceeds in rounds of bounded length, and the only requirement is that each player gets a chance to move in every round.

A subtle point about our results is that they refer to a *static* social information context, an assumption questionable in many settings. This is especially true for the results on the convergence rate of the ε-Nash dynamics, since during the convergence process, players can become aware of some initially unknown players. However, our results convey the message that the more the social information available, the better the situation becomes due to the fact that $\alpha(G)$ tends to decrease. So as the players collect social information and increase their social neighborhood, the ε-Nash dynamics keeps reaching better and better configurations. The entire process takes a polynomial number of steps, since there are $\mathrm{O}(n^2)$ acquaintances to be added to the social graph, and for each fixed social graph, the ε-Nash dynamics reaches a configuration with the desired properties in a polynomial number of steps.

Other Related Work. To the best of our knowledge, Gairing, Monien, and Tiemann [16] were the first to investigate the impact of incomplete social knowledge on the basic properties of weighted congestion games. They adopted a Bayesian approach, and mostly focused on parallel-link games.

Our information model can be regarded as a simplified version of the information model considered by Koutsoupias, Panagopoulou, and Spirakis [18]. Their model is based on a directed social graph, where each player knows the precise weights of the players in his social neighborhood, and only a probability distribution for the weights of the rest. Koutsoupias *et al.* obtained upper and lower bounds on the PoA for a very simple game with just two identical parallel links.

An alternative model for investigating the impact of social ignorance on the PoA for congestion games was suggested by Karakostas *et al.* [17]. In their model, a fraction of the players are totally ignorant to the presence of other players, and thus oblivious to the resource congestion when selecting their strategies, while the remaining players have full knowledge. Karakostas *et al.* considered non-atomic congestion games, and investigated how the PoA depends on the fraction of ignorant players.

After introducing the model of graphical congestion games in [7], Bilò *et al.* [8] considered the PoA and the PoS of graphical multicast cost sharing games, and proved that a central authority can dramatically decrease the PoA by enforcing a carefully selected social graph.

In an orthogonal approach, Ashlagi, Krysta, and Tennenholtz [4] associated the social graph not with the information available to the players, but with their individual cost. They suggested that the individual cost of each player is given by an *aggregation function* of the delays of the players in his social neighborhood. The aggregation function is also part of the social context, since it represents the players' attitude towards their neighbors.

2 Model and Preliminaries

For any integer $k \geq 1$, we denote $[k] \equiv \{1, \dots, k\}$. For a vector $x = (x_1, \dots, x_n)$, we denote $x_{-i} \equiv (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$ and $(x_{-i}, x'_i) \equiv (x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_n)$.

Weighted Congestion Games. A *congestion game* with weighted players is defined by a set $V = [n]$ of players, a positive integer[7] weight w_i associated with each player i, a set R of resources, a strategy space $\Sigma_i \subseteq 2^R \setminus \{\emptyset\}$ for each player i, and a non-decreasing delay function $d_e : \mathbb{N} \mapsto \mathbb{R}_{\geq 0}$ associated with each resource e. The game is (or the players are) unweighted if $w_i = 1$ for all $i \in [n]$. A congestion game has *symmetric strategies* if all players have a common strategy space. Throughout

[7] Throughout this paper, we restrict our attention to players with integral weights. This is essentially without loss of generality, since any game with fractional player weights can be turned into a game with integral player weights without affecting the properties considered in this paper.

this paper, we consider *linear* congestion games, where every resource e is associated with a linear delay function $d_e(x) = a_e x + b_e$, $a_e, b_e \geq 0$.

A *configuration* is a vector $s = (s_1, \ldots, s_n)$ consisting of a strategy $s_i \in \Sigma_i$ for each player i. For every resource e, we let $s_e = \sum_{i: e \in s_i} w_i$ denote the congestion induced on e by s. The (actual) cost of player i in the configuration s is $c_i(s) = \sum_{e \in s_i} w_i(a_e s_e + b_e)$.

Graphical Congestion Games. The new ingredient in the definition of graphical congestion games is the *social graph* $G(V, E)$, which is defined on the set of players V and contains an edge $\{i, j\} \in E$ between each pair of players i, j that know each other. In this work, we consider graphical congestion games weighted players and simple undirected social graphs, although one can define (and in certain cases, obtain similar results for) graphical games with directed social graphs.

Given a graphical congestion game with a social graph $G(V, E)$, let $\alpha(G)$ be the independence number of G, i.e. the cardinality of the largest set of players that do not know each other. For every configuration s and every resource e, let $V_e(s) = \{i \in V : e \in s_i\}$ be the set of players using e in s, let $G_e(s)(V_e(s), E_e(s))$ be the social subgraph of G induced by $V_e(s)$, and let $n_e(s) = |V_e(s)|$ and $m_e(s) = |E_e(s)|$. For each player i (not necessarily belonging to $V_e(s)$), let $\Gamma_e^i(s) = \{j \in V_e(s) : \{i, j\} \in E\}$ be the social neighborhood of player i among the players using e in s, and let $\deg_e^i(s) = |\Gamma_e^i(s)|$ be the number of i's neighbors using e in s.

In any configuration s, each player i is aware of a *presumed congestion* $s_e^i = w_i + \sum_{j \in \Gamma_e^i(s)} w_j$ on each resource e, and of his *presumed cost* $p_i(s) = \sum_{e \in s_i} w_i(a_e s_e^i + b_e)$. We observe that the presumed cost coincides with the actual cost if the social graph is complete.

For graphical congestion games, a configuration s is a *pure Nash equilibrium* (PNE) if no player can improve his *presumed cost* by unilaterally changing his strategy. Formally, s is a PNE if for every player i and all strategies $\sigma_i \in \Sigma_i$, $p_i(s) \leq p_i(s_{-i}, \sigma_i)$.

Social Cost, Price of Anarchy, and Price of Stability. In the first part of this work, we are interested in quantifying the inefficiency of PNE for graphical linear congestion games with weighted players. We evaluate configurations using the objective of *total actual cost*. The total cost $C(s)$ of a configuration s is the sum of players' actual costs in s, i.e. $C(s) = \sum_{i=1}^n c_i(s) = \sum_{e \in E}(a_e s_e^2 + b_e s_e)$. The *optimal configuration*, denoted o, minimizes the total cost among all configurations.

The (pure) *Price of Anarchy* (PoA) of a graphical congestion game $\mathcal{C}$ is the maximum ratio $C(s)/C(o)$ over all PNE s of $\mathcal{C}$. The (pure) *Price of Stability* (PoS) of $\mathcal{C}$ is the minimum ratio $C(s)/C(o)$ over all PNE s of $\mathcal{C}$. In other words, the PoA (resp. PoS) is equal to $C(s)/C(o)$, where s is $\mathcal{C}$'s PNE of maximum (resp. minimum) total cost.

Other Notions of Cost. We also consider the *total presumed cost* $P(s)$ of a configuration s, defined as $P(s) = \sum_{i=1}^n p_i(s)$. For unweighted graphical linear games, we observe that

$$P(s) = \sum_{e \in R} \sum_{i \in V_e(s)} (a_e(\deg_e^i + 1) + b_e) = \sum_{e \in R} (a_e(2m_e(s) + n_e(s)) + b_e n_e(s))$$

For any configuration s, $P(s) \leq C(s)$, which holds with equality if the social graph is complete. We let o' denote the configuration of minimum total presumed cost. The Price of Anarchy with respect to the total presumed cost is the maximum ratio $P(s)/P(o')$ over all PNE s.

Moreover, it is helpful to define the total singleton cost $U(s) = \sum_{i=1}^n \sum_{e \in s_i} w_i(a_e w_i + b_e)$. For any configuration s, $U(s) \leq P(s)$, which holds with equality if the social graph is an independent set.

Potential Functions. A function Φ that assigns a non-negative number $\Phi(s)$ to every configuration s is a *potential function* for a (graphical congestion) game if when a player i moves from his current strategy s_i to a new strategy $s_i' \in \Sigma_i$, the difference in the potential value equals the difference in the (presumed) cost of player i, i.e. $\Phi(s_{-i}, s'_i) - \Phi(s) = p_i(s_{-i}, s'_i) - p_i(s)$. If a game admits a potential function, its PNE correspond to the local minima of the potential function.

Best Responses, Improvement Moves, and Approximate Equilibria. A strategy $s'_i \in \Sigma_i$ is a *best response* of player i to a configuration s if for all strategies $\sigma_i \in \Sigma_i$, $p_i(s_{-i}, s'_i) \leq p_i(s_{-i}, \sigma_i)$. Given a configuration s, we let $\Delta(s) = \sum_{i=1}^{n}(p_i(s) - p_i(s_{-i}, s'_i))$ denote the sum of the improvements on the presumed cost if each player i moves from his current strategy to his best response s'_i.

A strategy $\sigma_i \in \Sigma_i$ is an *improvement move* of player i in a configuration s if $p_i(s_{-i}, \sigma_i) < p_i(s)$. Given an $\varepsilon \in (0, 1)$, a strategy $\sigma_i \in \Sigma_i$ is an (improvement) ε-*move* of player i in a configuration s if $p_i(s_{-i}, \sigma_i) < (1 - \varepsilon)p_i(s)$, i.e. if player i moving from his current strategy s_i to σ_i improves his presumed cost by a factor more than ε.

For a (graphical congestion) game that admits a potential function, every improvement move decreases the potential. Therefore, the Nash dynamics, i.e. any sequence of improvement moves, converges to a PNE in a finite number of steps. Similarly, the ε-Nash dynamics, i.e. any sequence of ε-moves, converges to a *pure Nash* ε-*equilibrium* (ε-PNE), i.e. a configuration where no player has an ε-move available. Formally, a configuration s is a ε-PNE if for every player i and all strategies $\sigma_i \in \Sigma_i$, $(1 - \varepsilon)c_i(s) \leq c_i(s_{-i}, \sigma_i)$.

In the second part of this work, we are interested in bounding the number of ε-moves so as to reach an approximately optimal configuration with approximation ratio arbitrarily close to the PoA of the game, and for unweighted graphical linear games with symmetric strategies, an ε-PNE.

3 Potential Function and Cost Approximation

Potential Function. We first show that graphical linear congestion games with weighted players admit a potential function, and thus any sequence of improvement moves converges to a PNE.

Theorem 1. *Every graphical linear congestion game with weighted players admits a potential function, and thus a pure Nash equilibrium.*

Proof. Let s be any configuration of a graphical linear congestion game $\mathcal{C}$ with weighted players. We show that

$$\Phi(s) = \sum_{e \in R} \left[a_e \left(\sum_{i \in V_e(s)} w_i^2 + \sum_{\{i,j\} \in E_e(s)} w_i w_j \right) + b_e \sum_{i \in V_e(s)} w_i \right] = \frac{P(s) + U(s)}{2}$$

is a potential function for $\mathcal{C}$. To this end, let i be a player switching from his strategy s_i in s to a strategy $s'_i \in \Sigma_i$, and let $s' = (s_{-i}, s'_i)$ be the resulting configuration. By simple algebra,

$$\begin{aligned} \Phi(s') - \Phi(s) &= \sum_{e \in s'_i \setminus s_i} \left[a_e \left(w_i^2 + \sum_{j \in \Gamma_e^i(s')} w_i w_j \right) + b_e w_i \right] - \sum_{e \in s_i \setminus s'_i} \left[a_e \left(w_i^2 + \sum_{j \in \Gamma_e^i(s)} w_i w_j \right) + b_e w_i \right] \\ &= w_i \sum_{e \in s'_i \setminus s_i} (a_e {s'}_e^i + b_e) - w_i \sum_{e \in s_i \setminus s'_i} (a_e s_e^i + b_e) \\ &= p_i(s') - p_i(s) \end{aligned}$$

Therefore, the difference in Φ due to player i switching from s_i to s'_i equals the corresponding difference in i's individual presumed cost. Consequently, Φ is an exact potential function for the graphical linear congestion game $\mathcal{C}$ with weighted players. □

Approximating Actual Cost with Presumed Cost. Our upper bounds on the PoA and the PoS are based on the fact that for any configuration s, the total actual cost $C(s)$ is an $\alpha(G)$-approximation of the total presumed cost $P(s)$, where $\alpha(G)$ is the *independence number* of the social graph G. Rather surprisingly, the worst-case ratio $C(s)/P(s)$ does not depend on the maximum weight set of players that do not know each other, it only depends on their maximum number.

Lemma 1. *Let $\mathcal{C}$ be any graphical linear congestion game with weighted players, let $G(V, E)$ be the social graph associated with $\mathcal{C}$, and let $\alpha(G)$ be the independence number of G. Then for any configuration s, $C(s) \leq \alpha(G)P(s)$.*

Proof. Let s be any configuration, let e be any resource, let $G_e(s)(V_e(s), E_e(s))$ be the subgraph induced by the players in $V_e(s)$, and let $C_e(s) = a_e s_e^2 + b_e s_e$ and $P_e(s) = \sum_{i \in V_e(s)} w_i(a_e s_e^i + b_e)$. We show that $C_e(s) \leq \alpha(G_e(s))P_e(s)$. The lemma follows since $\alpha(G_e(s)) \leq \alpha(G)$ for any configuration s and resource e, and since $C(s) = \sum_{e \in R} C_e(s)$ and $P_e(s) = \sum_{e \in R} P_e(s)$. For simplicity and since we consider a fixed configuration s throughout this proof, we omit the dependence of the social subgraph and its parameters on s.

To provide some intuition behind the proof, we consider a graph G_e with $n_e = k\alpha(G_e)$ vertices, where k is a positive integer, whose vertices can be partitioned into $\alpha(G_e)$ cliques of size k each. Let $s_e(\ell)$ be the total weight of the vertices in the ℓ-th clique, $\ell = 1, \ldots, \alpha(G_e)$. We observe that

$$P_e(s) \geq \sum_{\ell=1}^{\alpha(G_e)} (a_e s_e^2(\ell) + b_e s_e(\ell)) \geq a_e s_e^2/\alpha(G) + b_e s_e \geq C_e(s)/\alpha(G)\,,$$

where the penultimate inequality follows from the Cauchy-Schwarz inequality. In addition, (we below prove that) if G_e does not include any edges other than those in the cliques, then it has the minimum number of edges that a graph with so many vertices and such a maximum independent set can have. Intuitively, such a graph should maximize the ratio $C(s)/P(s)$.

For a formal proof, we first show that the inequality $C_e(s) \leq \alpha(G_e)P_e(s)$ is valid if all players have unit weights, and then we reduce the weighted case to the unweighted one. In the unweighted case, it suffices to establish a lower bound on the number of edges m_e, since for unweighted graphical games, $P_e(s)$ only depends on m_e, and not on which players are connected by an edge.

We now let $k = n_e/\alpha(G_e) \geq 1$ be some rational number, and let $r = k - \lfloor k \rfloor$ (resp. $k - r$) be k's fractional (resp. integral) part. We partition the vertices of G_e into a sequence of at least $\lceil k \rceil$ independent sets of non-increasing cardinality as follows: We start with the entire graph $G_e^{(1)} = G_e$ and $\ell = 1$. As long as $G_e^{(\ell)}$ is non-empty, we find a maximum independent set I_ℓ of $G_e^{(\ell)}$, obtain the graph $G_e^{(\ell+1)}$ by removing the vertices of I_ℓ from $G_e^{(\ell)}$, increase ℓ by one, and iterate. Let $q \geq \lceil k \rceil$ be the number of independent sets obtained by this decomposition.

The crucial observation is that since I_ℓ is a maximum independent set of $G_e^{(\ell)}$, for every $j = \ell + 1, \ldots, q$, each vertex $u \in I_j$ is connected by an edge to at least one vertex in I_ℓ. Otherwise, u could be added to I_ℓ, which results in an independent set larger than I_ℓ. Hence, for every $j = 2, \ldots, q$, each vertex $u \in I_j$ is incident to at least $j - 1$ edges connecting it to vertices in the independent sets $I_1, \ldots, I_{j-1}$. Therefore, the total number of edges in G_e is at least $\sum_{j=2}^{q}(j-1)|I_j|$. Since $\sum_{j=2}^{q} |I_j| = n_e - \alpha(G_e)$ and $|I_j| \leq \alpha(G_e)$ for all $j \in [q]$, the lower bound on the total number of edges is minimized when $q = \lceil k \rceil$, $|I_j| = \alpha(G_e)$ for all $j = \{1, \ldots, \lfloor k \rfloor\}$, and $|I_{\lceil k \rceil}| = r\alpha(G_e)$. Thus we obtain that

$$m_e \geq \frac{(k-r)(k-r-1)}{2}\alpha(G_e) + (k-r)r\alpha(G_e) = \frac{(k-r)(k+r-1)}{2}\alpha(G_e) \qquad (1)$$

Therefore,

$$\begin{aligned}\alpha(G_e)P_e(s) &= \alpha(G_e)[a_e(2m_e+n_e)+b_en_e] \\ &\geq a_e[(k-r)(k+r-1)+k]\alpha^2(G_e)+b_ek\alpha^2(G_e) \\ &\geq a_ek^2\alpha^2(G_e)+b_ek\alpha(G_e)=C_e(s)\,,\end{aligned}$$

where we use (1) for the first inequality, and that $r \geq r^2$, since $r \in [0,1)$, for the second inequality. This concludes the proof of the lemma for the unweighted case.

If the players have different weights w_i, $i \in V$, we create a new graph G'_e by replacing each player / vertex $i \in V_e$ with a clique Q_i of w_i vertices / players, each of unit weight. For each edge $\{i,j\} \in E_e$, we add w_iw_j edges to G'_e that connect every vertex in the clique Q_i to every vertex in the clique Q_j. Thus G'_e has $\sum_{i\in V_e} w_i$ vertices and $\sum_{i\in V_e} w_i(w_i-1)/2+\sum_{\{i,j\}\in E_e} w_iw_j$ edges.

We claim that (i) the transformation above does not affect $P_e(s)$ and $C_e(s)$, and that (ii) $\alpha(G_e) = \alpha(G'_e)$. The former claim follows directly from the definitions of G'_e, $C_e(s)$, and $P_e(s)$. The latter claim holds because at most one vertex from each clique Q_i can be included in an independent set of G'_e, and two vertices from the cliques Q_i and Q_j are not connected by an edge in G'_e iff vertices i and j are not connected by an edge in G_e. Therefore, there is a correspondence between independent sets of a given size in G_e and G'_e, which implies the claim.

Since we have already established the lemma in the unweighted setting, $P_e(s) \leq \alpha(G'_e)C_e(s)$, which implies that $P_e(s) \leq \alpha(G_e)C_e(s)$ is valid even if the players have different weights. □

4 The Price of Anarchy and the Price of Stability

An Upper Bound on the Presumed Cost. We start by establishing an upper bound on the presumed cost $P(s)$ of any configuration s in terms of the optimal total cost $C(o)$ and the sum of improvements on the presumed cost $\Delta(s) = \sum_{i=1}^{n}(p_i(s) - p_i(s_{-i}, s'_i))$ if each player i switches from s_i to his best response s'_i. This upper bound is useful both in bounding the PoA and in establishing the fast convergence of the ε-Nash dynamics to approximately optimal configurations.

Lemma 2. *Let $\mathcal{C}$ be a graphical linear congestion game with weighted players arranged in a social graph G, let $\alpha(G)$ be the independence number of G, and let o be the optimal configuration. Then, for any configuration s,*

$$P(s) \leq \frac{\alpha(G)+2+\sqrt{\alpha^2(G)+4\alpha(G)}}{2}C(o)+2\Delta(s) \tag{2}$$

Proof. We follow the general approach of [6, Lemma 4.1] and [5, Theorem 3.1]. The difference is that we have to bound the total presumed cost in terms of the optimal total cost.

The presumed cost of each players i if he switches to his best response strategy s'_i is at most his presumed cost if he switches to his optimal strategy o_i. Thus,

$$p_i(s_{-i}, s'_i) \leq p_i(s_{-i}, o_i) \leq \sum_{e\in o_i} w_i(a_e(s_e+w_i)+b_e) \tag{3}$$

Summing up over all players, we obtain the following inequality, whose proof can be found in the Appendix, Section A.1:

$$\sum_{i=1}^{n} p_i(s_{-i}, s'_i) \leq \sqrt{C(o)\alpha(G)P(s)}+C(o) \tag{4}$$

By adding $\Delta(s)$ to both sides of (4) and dividing by $C(o)$, we obtain that:

$$\frac{P(s)}{C(o)} \leq \sqrt{\alpha(G)}\sqrt{\frac{P(s)}{C(o)}} + 1 + \frac{\Delta(s)}{C(o)} \tag{5}$$

Setting $\beta = P(s)/C(o)$ and $\gamma = \Delta(s)/C(o)$, (5) becomes $\beta^2 \leq \sqrt{\alpha(G)}\beta + 1 + \gamma$. By simple algebra, we obtain that $\beta \leq (\sqrt{\alpha(G)} + \sqrt{\alpha(G) + 4(1+\gamma)})/2$, which implies that

$$\frac{P(s)}{C(o)} \leq \frac{\alpha(G) + 2 + \sqrt{\alpha^2(G) + 4\alpha(G)}}{2} + 2\,\frac{\Delta(s)}{C(o)} \quad \square$$

The Price of Anarchy. Lemma 1 and Lemma 2 immediately imply the following upper bound on the PoA of graphical linear congestion games with weighted players.

Theorem 2. *For graphical linear congestion games with weighted players, the Price of Anarchy is at most $\alpha(G)(\alpha(G) + 2 + \sqrt{\alpha^2(G) + 4\alpha(G)})/2$, where $\alpha(G)$ denotes the independence number of the social graph G.*

Theorem 2 implies that the PoA of any graphical linear game with weighted players is less than $\alpha(G)(\alpha(G) + 2)$. Bilò *et al.* [7, Theorem 13] present a family of unweighted graphical games with n players arranged in a bipartite social graph G with $\alpha(G) = n/2$, for which the PoA is $\omega(\alpha(G))$ and can be as large as $\alpha^2(G)$. Next we present a different family of unweighted graphical games for which the ratio $\alpha(G)/n$ can take any value in $(0, \frac{1}{2}]$ and the PoA is at least $\alpha(G)(\alpha(G) + 1)$. Thus we show that as long as $\alpha(G) \leq n/2$, the upper bound of Theorem 2 is essentially tight.

Theorem 3. *For any integers $\ell \geq 1$ and $n \geq 2\ell$, there is a graphical linear congestion game with n unweighted players arranged in a social graph G with $\alpha(G) = \ell$, for which the PoA is $\ell(\ell + 1)$.*

Proof. We first show how to construct such a graphical game if $k = n/\ell$ is an integer. The social graph G is the complete k-partite graph with each independent set consisting of ℓ vertices. The resource set R consists of $k(\ell+1)$ resources e_i^j, $j \in [k]$, $i \in \{0\} \cup [\ell]$. All resources have delay function $d(x) = x$.

Each player has two strategies, the "short" one and the "long" one. For the i-th player in the j-th independent set, the "short" strategy is $\{e_i^j\}$, and the "long" strategy is to use all resources of the next independent set, i.e. $\{e_0^{(j \bmod k)+1}, e_1^{(j \bmod k)+1}, \ldots, e_\ell^{(j \bmod k)+1}\}$ (see also Fig. 1 in the Appendix).

The optimal configuration o assigns all players to their "short" strategies and achieves a total cost of $C(o) = k\ell$. On the other hand, there is a PNE s where all players use their "long" strategies. The presumed cost of each player in s is $\ell + 1$, equal to his presumed cost if he switches to his "short" strategy. The actual cost of each player in s is $\ell(\ell + 1)$, the total cost of s is $C(s) = k\ell^2(\ell + 1)$, and the PoA is $\ell(\ell + 1)$.

We can handle the case where n/ℓ is not an integer by letting $k = \lfloor n/\ell \rfloor$ and adding $n - k\ell < \ell$ "dummy" players. The "dummy" players are connected to all other players, and their only strategy is a "dummy" resource e with delay $d_e(x) = 0$. $\square$

Lemma 1 and Lemma 2 also imply the same upper bound on the PoA of graphical linear congestion games with weighted players if the PoA is calculated with respect to the total presumed cost.

Theorem 4. *For a graphical congestion game with weighted players, let o' be the configuration minimizing the total presumed cost, and let s be any PNE. Then,*

$$P(s) \leq \frac{\alpha(G)}{2}(\alpha(G) + 2 + \sqrt{\alpha^2(G) + 4\alpha(G)})P(o'),$$

where $\alpha(G)$ denotes the independence number of the social graph G.

Moreover, we show that the bound of Theorem 4 is essentially tight for social graphs G with $\alpha(G) \leq \sqrt{n/2}$, even for unweighted players. We emphasize that such a lower bound is best possible, since the Price of Anarchy with respect to the total presumed cost is at most n (see e.g. [7, Theorem 9], which can be easily generalized to the weighted case).

Theorem 5. *For any integers $\ell \geq 1$ and $n \geq 2\ell^2$, there is a graphical linear congestion game with n unweighted players arranged in a social graph G with $\alpha(G) = \ell$, for which the Price of Anarchy with respect to the total presumed cost is ℓ^2.*

Proof. We restrict our attention to the most interesting case where $k = n/\ell^2$ is an integer. We can easily handle the case where n/ℓ^2 is not an integer as in the proof of Theorem 3.

For any integer $k \geq 2$, we construct a graphical game with $k\ell^2$ unit weight players and PoA with respect to the total presumed cost equal to ℓ^2. The social graph G consists of k groups G_i, $i \in [k]$, where each group G_i consists of ℓ disjoint independent sets I_i^j, $j \in [\ell]$, with ℓ vertices each. The vertices within each group G_i are also partitioned into ℓ cliques of cardinality ℓ, with each clique including one vertex from each independent set I_i^j, $j \in [\ell]$. The edges between the vertices in the same clique are the only edges between vertices in the same group. Furthermore, all pairs of vertices from different groups are connected to each other by an edge (see also Fig. 2 in the Appendix). The resource set R consists of $k\ell$ resources e_i^j, $i \in [k]$, $j \in [\ell]$, one for each independent set I_i^j. All resources have delay function $d(x) = x$.

Each player has two strategies, the "short" one and the "long" one. For each player in the independent set I_i^j, the "short" strategy is $\{e_i^j\}$, and the "long" strategy is $\{e^1_{(i \bmod k)+1}, \ldots, e^\ell_{(i \bmod k)+1}\}$ (see also Fig. 3 in the Appendix).

The configuration o' that minimizes the total presumed cost assigns each player to his "short" strategy. Hence each player has presumed cost 1, and the total presumed cost is $P(o') = k\ell^2$. On the other hand, there is a PNE s where all players use their "long" strategies. The presumed cost of each player in s is ℓ^2, since all players from the same clique in their group are assigned to the same resources. If a player switches to his "short" strategy, his presumed cost is greater than ℓ^2, because all ℓ^2 players from the previous independent set are assigned to the corresponding resource. Hence s is indeed a PNE with total presumed cost $P(s) = k\ell^4$. Therefore, the PoA with respect to the total presumed cost is $P(s)/P(o') = \ell^2$. □

The Price of Stability. An upper bound on the PoS follows easily by the potential function of Theorem 1 and the bound of Lemma 2.

Theorem 6. *For graphical linear congestion games with weighted players, the Price of Stability is at most $\frac{2n\alpha(G)}{n+\alpha(G)}$, where n denotes the number of players and $\alpha(G)$ denotes the cardinality of the maximum independent set of the social graph G.*

Proof. Let $\mathcal{C}$ be a graphical linear congestion game with n weighted players and social graph G, let o be the optimal configuration, and let s be the configuration that minimizes the potential function Φ. Since s is a PNE of $\mathcal{C}$, it suffices to show that $C(s)/C(o) \leq \frac{2n\alpha(G)}{n+\alpha(G)}$. To this end, we observe that:

$$P(s)/2 + U(s)/2 = \Phi(s) \leq \Phi(o) \leq C(o)\,,$$

where the first inequality follows from the fact that s is a minimizer of Φ.

By Lemma 1, $P(s) \geq C(s)/\alpha(G)$. In addition,

$$U(s) = \sum_{e \in R} \left(a_e \textstyle\sum_{i \in V_e(s)} w_i^2 + b_e s_e \right) \geq \sum_{e \in R} (a_e s_e^2/n + b_e s_e) \geq C(s)/n\,,$$

where we use the Cauchy-Schwarz inequality. Therefore, $C(s)(\frac{1}{2\alpha(G)} + \frac{1}{2n}) \leq C(o)$. □

The following theorem presents a family of unweighted graphical linear games for which the PoS is almost as large as the independence number of the social graph, thus showing that the upper bound of Theorem 6 is essentially tight. Its proof can be found in the Appendix, Section A.2.

Theorem 7. *For any positive integers ℓ and $n \geq \ell$, and any $\varepsilon > 0$, there is a graphical linear congestion game with n unweighted players arranged in a social graph G with $\alpha(G) = \ell$, for which the Price of Stability is $\ell - \varepsilon$.*

5 Convergence Rate of the ε-Nash Dynamics

Convergence to Near Optimal Configurations. Employing the techniques of Awerbuch *et al.* [6], we show that the *largest improvement* ε-Nash dynamics reaches an approximately optimal configuration in a polynomial number of steps, where the approximation ratio is arbitrarily close to the PoA of the graphical congestion game. In each step of the largest improvement ε-Nash dynamics, among all players with an ε-move available, the player with the largest improvement on his presumed cost moves. The proof of the following theorem can be found in the Appendix, Section A.3.

Theorem 8. *Let $\mathcal{C}$ be a graphical linear congestion game with n weighted players arranged in a social graph $G(V, E)$, let s^* be a minimizer of the potential function Φ, and let $\frac{1}{8} \geq \delta \geq \varepsilon > 0$. Starting from an initial configuration s_0, the largest improvement ε-Nash dynamics reaches a configuration s with $C(s) \leq \frac{\alpha(G)}{2}(\alpha(G) + 2 + \sqrt{\alpha^2(G) + 4\alpha(G)})(1 + 8\delta)C(o)$ in $\mathrm{O}(\frac{n}{\delta} \log \frac{\Phi(s_0)}{\Phi(s^*)})$ steps.*

Furthermore, following the approach of [6, Theorem 3.3], we can establish a similar convergence time for the unrestricted ε-Nash dynamics, which proceeds in rounds of bounded length, and the only requirement is that each player gets a chance to move in every round. Due to lack of space, we defer the technical details to the full version of the paper.

Convergence to Approximate Equilibria. For graphical linear games with unweighted players and symmetric strategies, we employ the techniques of Chien and Sinclair [10] and show the stronger result that the largest improvement ε-Nash dynamics converges an ε-PNE in a polynomial number of steps. Compared to the bound implied by [10, Theorem 3.1] for symmetric linear congestion games, the convergence time increases by a factor up to n due to the players having different social neighborhoods. The proof of the following theorem can be found in the Appendix, Section A.4.

Theorem 9. *Let $\mathcal{C}$ be a graphical linear congestion game with symmetric strategies and n unweighted players, let s^* be a minimizer of the potential function Φ, and let $\varepsilon \in (1, 0)$. Starting from an initial configuration s_0, the largest improvement ε-Nash dynamics converges in $\mathrm{O}(\frac{n^2}{\varepsilon} \log \frac{\Phi(s_0)}{\Phi(s^*)})$ steps.*

Moreover, following the approach of [10, Theorem 4.1], we can establish a similar convergence time to an ε-PNE for the unrestricted ε-Nash dynamics, where the only requirement is that each player gets a chance to move in every round. Due to lack of space, we defer the technical details to the full version of the paper.

Acknowledgements: The first author thanks Piotr Krysta for many discussions on the effects of the social context on non-cooperative games, which discussions greatly contributed to shaping his view on the topic of this work.

References

1. H. Ackermann, H. Röglin, and B. Vöcking. On the Impact of Combinatorial Structre on Congestion Games. *Journal of the ACM*, 55(6), 2008.
2. S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact Price of Anarchy for Polynomial Congestion Games. In *Proc. of the 23st Symp. on Theoretical Aspects of Computer Science (STACS '06)*, LNCS 3884, pp. 218–229, 2006.
3. E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The Price of Stability for Network Design with Fair Cost Allocation. In *Proc. of the 45th IEEE Symp. on Foundations of Computer Science (FOCS '04)*, pp. 295–304, 2004.
4. I. Ashlagi, P. Krysta, and M. Tennenholtz. Social Context Games. In *Proc. of the 4th Workshop on Internet and Network Economics (WINE '08)*, LNCS 5385, pp. 675–683, 2008.
5. B. Awerbuch, Y. Azar, and A. Epstein. The Price of Routing Unsplittable Flow. In *Proc. of the 37th ACM Symp. on Theory of Computing (STOC '05)*, pp. 57–66, 2005.
6. B. Awerbuch, Y. Azar, A. Epstein, V. Mirrokni, and A. Skopalik. Fast Convergence to Nearly Optimal Solutions in Potential Games. In *Proc. of the 9th ACM Conf. on Electronic Commerce (EC '08)*, pp. 264–273, 2008.
7. V. Bilò, A. Fanelli, M. Flammini, and L. Moscardelli. Graphical Congestion Games. In *Proc. of the 4th Workshop on Internet and Network Economics (WINE '08)*, LNCS 5385, pp. 70–81, 2008.
8. V. Bilò, A. Fanelli, M. Flammini, and L. Moscardelli. When Ignorance Helps: Graphical Multicast Cost Sharing Games. In *Proc. of the 33rd Symp. on Mathematical Foundations of Computer Science (MFCS '08)*, LNCS 5162, pp. 108–199, 2008.
9. I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, and L. Moscardelli. Tight Bounds for Selfish and Greedy Load Balancing. In *Proc. of the 33th Colloq. on Automata, Languages and Programming (ICALP '06)*, LNCS 4051, pp. 311–322, 2006.
10. S. Chien and A. Sinclair. Convergece to Approximate Nash Equilibria in Congestion Games. In *Proc. of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA '07)*, pp. 169–178, 2007.
11. G. Christodoulou and E. Koutsoupias. On the Price of Anarchy and Stability of Correlated Equilibria of Linear Congestion Games. In *Proc. of the 13th European Symposium on Algorithms (ESA '05)*, LNCS 3669, pp. 59–70, 2005.
12. G. Christodoulou and E. Koutsoupias. The Price of Anarchy of Finite Congestion Games. In *Proc. of the 37th ACM Symp. on Theory of Computing (STOC '05)*, pp. 67–73, 2005.
13. G. Christodoulou, E. Koutsoupias, and P. Spirakis. On the Performance of Approximate Equilibria in Congestion Games. In *Proc. of the 17th European Symposium on Algorithms (ESA '09)*, to appear, 2009.
14. A. Fabrikant, C. Papadimitriou, and K. Talwar. The Complexity of Pure Nash Equilibria. In *Proc. of the 36th ACM Symp. on Theory of Computing (STOC '04)*, pp. 604–612, 2004.
15. D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish Unsplittable Flows. *Theoretical Computer Science*, 348, pp. 226–239, 2005.
16. M. Gairing, B. Monien, and K. Tiemann. Selfish Routing with Incomplete Information. *Theory of Computing Systems*, 42, pp. 91–130, 2008.
17. G. Karakostas, T. Kim, A. Viglas, and H. Xia. Selfish Routing with Oblivious Users. In *14th Colloq. on Structural Information and Communication Complexity (SIROCCO '07)*, LNCS 4474, pp. 318–327, 2007.
18. E. Koutsoupias, P. Panagopoulou, and P. Spirakis. Selfish Load Balancing Under Partial Knowledge. In *Proc. of the 32nd Symp. on Mathematical Foundations of Computer Science (MFCS '07)*, LNCS 4708, pp. 609–620, 2007.
19. E. Koutsoupias and C. Papadimitriou. Worst-case Equilibria. In *Proc. of the 16th Symp. on Theoretical Aspects of Computer Science (STACS '99)*, LNCS 1563, pp. 404–413, 1999.
20. R.W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2, pp. 65–67, 1973.
21. A. Skopalik and B. Vöcking. Inapproximability of Pure Nash Equilibria. In *Proc. of the 40th ACM Symp. on Theory of Computing (STOC '08)*, pp. 355–364, 2008.

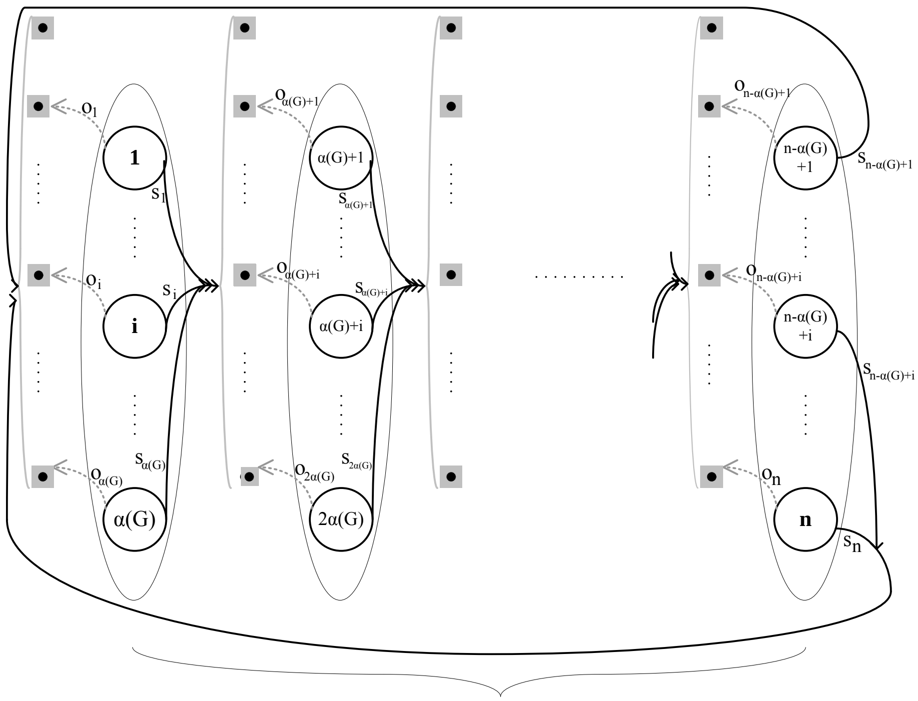

Fig. 1. The players' strategy space, the optimal configuration o, and the PNE configuration s in the proof of Theorem 3. The grey dotted arrows represent the assignment of each player in o, while the black solid arrows represent the assignment of each player in s.

A Appendix

A.1 The Proof of (4)

Summing up (3) over all players, we obtain that:

$$\begin{aligned}\sum_{i=1}^{n} p_i(s_{-i}, s'_i) &\le \sum_{i=1}^{n}\sum_{e\in o_i}(a_e s_e w_i + a_e w_i^2 + b_e w_i)\\ &\le \sum_{e\in R} a_e s_e o_e + \sum_{e\in R}(a_e o_e^2 + b_e o_e)\\ &\le \sqrt{\sum_{e\in R}(a_e o_e^2 + b_e o_e)\sum_{e\in R}(a_e s_e^2 + b_e s_e)} + C(o)\\ &= \sqrt{C(o)C(s)} + C(o)\\ &\le \sqrt{C(o)\alpha(G)P(s)} + C(o)\end{aligned}$$

For the second inequality, we rearrange the sums and use that $\sum_{i:e\in o_i} w_i = o_e$ and that $\sum_{i:e\in o_i} w_i^2 \le o_e^2$. The third inequality follows from the Cauchy-Schwarz inequality (see also [5, Lemma 3.1]). The last inequality follows from Lemma 1. □

A.2 The Proof of Theorem 7

We restrict our attention to the case where $k = n/\ell$ is an integer. We can easily handle the case where n/ℓ is not an integer as in the proof of Theorem 3.

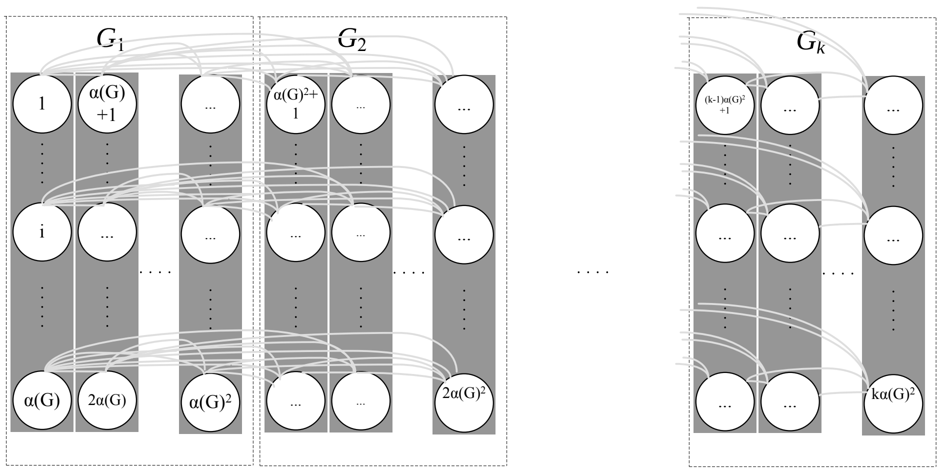

Fig. 2. The social graph in the proof of Theorem 5.

For any integer $k \geq 1$, we construct a graphical game with $k\ell$ unit weight players and symmetric strategies. The social graph G is the complete k-partite graph whose independent sets consist of ℓ vertices each. The resource set R consists of $k\ell$ resources e_j, $j \in [k\ell]$, and the strategy space of each player consists of all singleton subsets of R (i.e. the strategy space is determined by $k\ell$ parallel links). Let $\delta_1, \ldots, \delta_{k\ell}$ be appropriately small positive numbers such that $0 < \delta_1 < \cdots < \delta_{k\ell}$ (e.g. it suffices to set $\delta_j = j/n^\beta$ with $\beta \geq 3$), and let $\delta = \sum_{j=1}^{k\ell} \delta_j$. The delays are $d_{e_j}(x) = x + \delta_j$, for $j \in [k\ell]$.

The optimal configuration o assigns one player to each resource and achieves a total cost of $C(o) = k\ell + \delta$. In any PNE s, players from different independent sets use different resources, and all players in the same independent set use one of the first k resources. Therefore, the total cost of any PNE s is $C(s) \geq k\ell^2$, and the PoS is at least $\ell - \varepsilon$, provided that δ_i's are selected so that $\delta < \varepsilon/k$. □

A.3 The Proof of Theorem 8

For simplicity of notation, we let $\rho = (\alpha(G) + 2 + \sqrt{\alpha^2(G) + 4\alpha(G)})/2$ denote the multiplier of $C(o)$ in (2). We show that as long as the current configuration s has $P(s) > \rho(1 + 8\delta)C(o)$, there is a player with an ε-move that decreases the potential function by at least $\frac{\delta}{n}\Phi(s)$. Since the potential function decreases by a factor of $1 - \frac{\delta}{n}$ in each step, after $\mathrm{O}(\frac{n}{\delta} \log \frac{\Phi(s_0)}{\Phi(s^*)})$ steps a configuration s with $P(s) \leq \rho(1 + 4\delta)C(o)$ is reached. Then Lemma 1 implies that $C(s) \leq \alpha(G)\rho(1 + 8\delta)C(o)$.

More precisely, if the current configuration s has $P(s) > \rho(1 + 8\delta)C(o)$, it is not and ε-PNE. Otherwise, $\Delta(s) \leq \varepsilon P(s)$, and (2) implies that $P(s) < \rho(1 + 8\delta)C(o)$, since $\delta \in [\frac{1}{8}, \varepsilon]$. Thus we let $U \neq \emptyset$ be the set of players with an ε-move available in s, and let $\Delta_U(s) = \sum_{i \in U}(p_i(s) - p_i(s_{-i}, s'_i))$, where s'_i denotes the best response of player i to s. Then $\Delta_U(s) > \delta P(s)$, since otherwise $\Delta(s) \leq (\varepsilon + \delta)P(s)$, and (2) implies that $P(s) < \rho(1 + 8\delta)C(o)$. Therefore, there is a player $i \in U$ for whom moving from his current strategy s_i to his best response s'_i is an ε-move, and $p_i(s) - p_i(s_{i-1}, s'_i) > \frac{\delta}{n}\Phi(s)$, i.e. i moving from s_i to s'_i decreases the potential function by a factor greater than $1 - \frac{\delta}{n}$. □

A.4 The Proof of Theorem 9

We show that as long as the current configuration s is not an ε-PNE, there is a player i with an ε-move that decreases the potential function by at least $\varepsilon\Phi(s)/n^2$. Since the potential function decreases by a factor of $1 - \frac{\varepsilon}{n^2}$ in each step, after $\mathrm{O}(\frac{n^2}{\varepsilon} \log \frac{\Phi(s_0)}{\Phi(s^*)})$ steps an ε-PNE is reached.

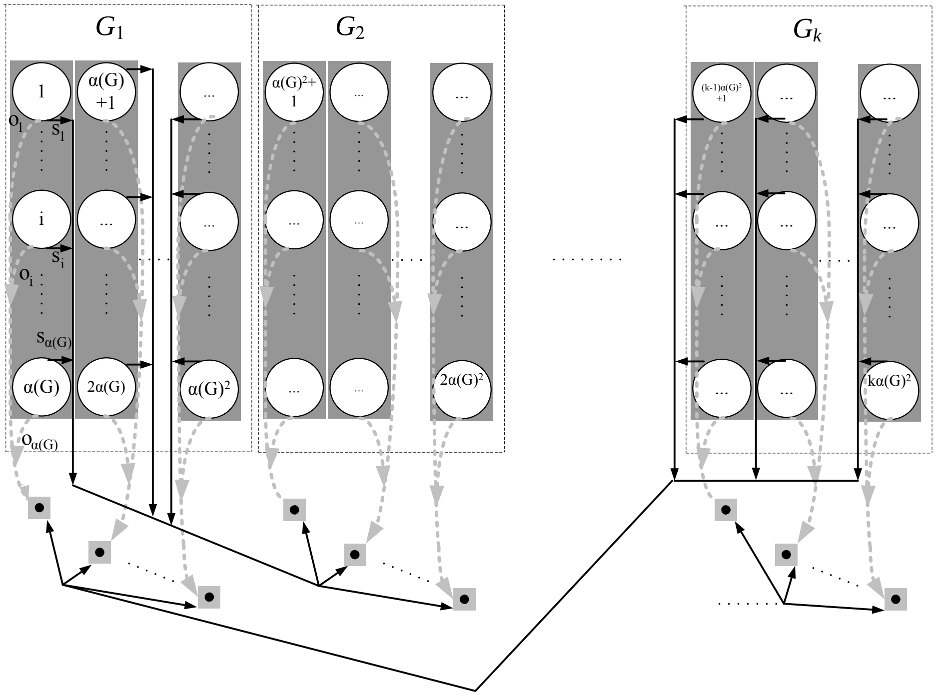

Fig. 3. The players' strategy space, the configuration o' of minimum total presumed cost, and the PNE configuration s in the proof of Theorem 5. The grey dotted arrows represent the assignment of each player in o', while the black solid arrows represent the assignment of each player in the s.

More precisely, let i be the next player to move from his strategy s_i in the current configuration s to his best response s'_i, and let j be the player with the largest presumed cost in s. We observe that $p_j(s) \geq \Phi(s)/n$.

If player j has also an ε-move in s, the decrease in his presumed cost is at least $\varepsilon p_j(s)$. Since player i is the one who moves, the decrease in i's presumed cost (and in the potential function) is at least $\varepsilon p_j(s) \geq \varepsilon\Phi(s)/n$.

If player j does not have an ε-move in s, we show that $p_i(s) > p_j(s)/n$. This is sufficient, since the decrease in i's presumed cost (and in the potential function) when player i moves from s_i to s'_i is at least $\varepsilon\Phi(s)/n^2$. To prove that $p_i(s) > p_j(s)/n$, we observe that, since all players have a common set of strategies, player j can move to i's best response s'_i and have a presumed cost:

$$p_j(s_{-j}, s'_i) \leq \sum_{e \in s'_i} (a_e(s^j_e + 1) + b_e) \leq n \sum_{e \in s'_i} (a_e + b_e) \leq n p_i(s_{-i}, s'_i)$$

Since player i has an ε-move in s but player j does not,

$$(1-\varepsilon)p_i(s) > p_i(s_{-i}, s'_i) \geq p_j(s_{-j}, s'_i)/n \geq (1-\varepsilon)p_j(s)/n$$

Therefore, $p_i(s) > p_j(s)/n$, which concludes the proof. □