

Atomic Selfish Routing in Networks: A Survey*

Spyros Kontogiannis^{†,‡}

Paul Spirakis[†]

May 8, 2005

Abstract

In this survey we present some recent advances in the *atomic congestion games* literature. Our main focus is on a special case of congestion games, called *network congestion games*, which is of particular interest for the networking community. The algorithmic questions that we are interested in have to do with the *existence* of pure Nash equilibria, the *efficiency* of their construction when they exist, as well as the *gap* of the best/worst (mixed in general) Nash equilibria from the social optima in such games, typically called the *Price of Anarchy* and the *Price of Stability* respectively.

1 Introduction

Consider a model where selfish individuals (henceforth called **players**) in a communication network having varying service demands compete for some shared resources. The quality of service provided by a resource decreases with its *congestion*, ie, the amount of demands of the players willing to be served by it. Each player may reveal its actual, unique choice of a subset of resources (called a *pure strategy*) that satisfies his service demand, or he may reveal a probability distribution for choosing (independently of other players' choices) one of the possible (satisfactory for him) subsets of resources (called a *mixed strategy*). The players determine their actual behavior based on other players' behaviors, but they do not cooperate. We are interested in situations where the players have reached some kind of stable state, ie, an equilibrium. The most popular notion of equilibrium in non-cooperative game theory is the *Nash equilibrium*: a “stable point” among the players, from which no player is willing to deviate unilaterally. In [23], the notion of the *coordination ratio* or *price of anarchy* was introduced as a means for measuring the performance degradation due to lack of players' coordination when sharing common goods. A more recent measure of performance is the *price of stability* [2], capturing the gap between the best possible Nash Equilibrium and the globally optimal solution. This measure is crucial for the network designer's perspective, who would like to propose (rather than let the players end up in) a Nash equilibrium (from which no player would like to defect unilaterally) that is as close to the optimum as possible.

A realistic scenario for the above model is when *unsplittable* traffic demands are routed selfishly in general networks with load-dependent edge delays. When the underlying network

*This work was partially supported by the EU within 6th Framework Programme under contract 001907 (DELIS).

[†] Research Academic Computer Technology Institute, Riga Feraiou 61, 26222 Patra, Greece. {kontog,spirakis}@cti.gr

[‡] Computer Science Department, University of Ioannina, 45110 Ioannina, Greece. kontog@cs.uoi.gr

consists of two nodes and parallel links between them, there has been an extensive study on the existence and computability of equilibria, as well as on the price of anarchy. In this survey we study the recent advances in the more general case of arbitrary *congestion games*. When the players have identical traffic demands, the congestion game is indeed isomorphic to an *exact potential game* ([29], see also theorem 1 of this survey) and thus always possesses a *pure Nash equilibrium*, ie, an equilibrium where each player adopts a pure strategy. We shall see that varying demands of the players crucially affect the nature of these games, which are no longer isomorphic to exact potential games. We also present some results in a variant of congestion games, where the players' payoffs are not resource-dependent (as is typically the case in congestion games) but *player-specific*.

1.1 Roadmap

In section 2 we formally define the congestion games and their variants considered in this survey. We also give some game-theoretic definitions. In section 3 we present most of the related work in the literature, before presenting in detail some of the most significant advances in the area. In section 4 we present some of the most important results concerning unweighted congestion games and their connection to the *potential games*. In section 5 we study some complexity issues of unweighted congestion games. In section 6 we presents Milchtaich's extension of congestion games to allow *player-specific* payoffs, whereas in section 5.2 we study some existence and computability issues of PNE in weighted congestion games. Finally, in section 7 we study the price of anarchy of weighted congestion games. We close this survey with some concluding remarks and unresolved questions.

2 The Model

Consider having a set of resources E in a system. For each $e \in E$, let $d_e(\cdot)$ be the **delay** per player that requests her service, as a function of the total usage (ie, the *congestion*) of this resource by all the players. Each such function is considered to be *non-decreasing* in the total usage of the corresponding resource. Each resource may be represented by a pair of points: an entry point to the resource and an exit point from it. So, we represent each resource by an arc from its entry point to its exit point and we associate with this arc the **charging cost** (eg, the delay as a function of the load of this resource) that each player has to pay if he is served by this resource. The entry/exit points of the resources need not be unique; they may coincide in order to express the possibility of offering a *joint service* to players, that consists of a sequence of resources. We denote by V the set of all entry/exit points of the resources in the system. Any nonempty collection of resources corresponding to a directed path in $G \equiv (V, E)$ comprises an **action** in the system.

Let $N \equiv [n]^1$ be the set of players, each willing to adopt some action in the system. $\forall i \in N$, let w_i denote player i 's **traffic demand** (eg, the flow rate from a source node to a destination node), while $\mathcal{P}^i \equiv \{a_1^i, \dots, a_{m_i}^i\} \subseteq 2^E \setminus \emptyset$ (for some $m_i \geq 2$) is the collection of actions, any of which would satisfy player i (eg, alternative routes from a source to a destination node, if G represents a communication network). The collection \mathcal{P}^i is called the *action set* of player i and each of its elements contains at least one resource. Any n -tuple $\varpi \in \mathcal{P} \equiv \times_{i=1}^n \mathcal{P}^i$ is a **pure strategies profile**, or a **configuration** of the players. Any

¹ $\forall k \in \mathbb{N}$, $[k] \equiv \{1, 2, \dots, k\}$.

real vector $\mathbf{p} = (\mathbf{p}^1, \mathbf{p}^2, \dots, \mathbf{p}^n)$ s.t. $\forall i \in N, \mathbf{p}^i \in \Delta(\mathcal{P}^i) \equiv \{\mathbf{z} \in [0, 1]^{m_i} : \sum_{k=1}^{m_i} z_k = 1\}$ is a probability distribution over the set of allowable actions for player i , is called a **mixed strategies profile** for the n players.

A **congestion model** $((\mathcal{P}^i)_{i \in N}, (d_e)_{e \in E})$ typically deals with players of identical demands, and thus the resource delay functions depend only on the *number* of players adopting each action ([31, 29, 12]). In the more general case, ie, a **weighted congestion model** is the tuple $((w_i)_{i \in N}, (\mathcal{P}^i)_{i \in N}, (d_e)_{e \in E})$. That is, we allow the players to have different (but fixed) demands for service (denoted by their weights) from the whole system, and thus affect the resource delay functions in a different way, depending on their own weights. The *weighted congestion game* $\Gamma \equiv (N, E, (w_i)_{i \in N}, E, (\mathcal{P}^i)_{i \in N}, (d_e)_{e \in E})$ associated with this model, is the game in strategic form with the set of players N and players' demands $(w_i)_{i \in N}$, the set of shared resources E , the action sets $(\mathcal{P}^i)_{i \in N}$ and players' cost functions $(\lambda_{\varpi^i}^i)_{i \in N, \varpi^i \in \mathcal{P}^i}$ defined as follows: For any configuration $\varpi \in \mathcal{P}$ and $\forall e \in E$, let $\Lambda_e(\varpi) = \{i \in N : e \in \varpi^i\}$ be the set of players wishing to exploit resource e according to ϖ (called the **view** of resource e wrt configuration ϖ). We also denote by $x_e(\varpi) \equiv |\Lambda_e(\varpi)|$ the *number* of players using resource e wrt ϖ , whereas $\theta_e(\varpi) \equiv \sum_{i \in \Lambda_e(\varpi)} w_i$ is the **load** of e wrt to ϖ . The **cost** $\lambda^i(\varpi)$ **of player** i **for adopting strategy** $\varpi^i \in \mathcal{P}^i$ in a given configuration ϖ is equal to the cumulative delay $\lambda_{\varpi^i}(\varpi)$ of all the resources comprising this action:

$$\lambda^i(\varpi) = \lambda_{\varpi^i}(\varpi) = \sum_{e \in \varpi^i} d_e(\theta_e(\varpi)). \quad (1)$$

On the other hand, for a mixed strategies profile \mathbf{p} , the **(expected) cost of player** i **for adopting strategy** $\varpi^i \in \mathcal{P}^i$ wrt \mathbf{p} is

$$\lambda_{\varpi^i}^i(\mathbf{p}) = \sum_{\varpi^{-i} \in \mathcal{P}^{-i}} P(\mathbf{p}^{-i}, \varpi^{-i}) \cdot \sum_{e \in \varpi^i} d_e(\theta_e(\varpi^{-i} \oplus \varpi^i)) \quad (2)$$

where, $\varpi^{-i} \in \mathcal{P}^{-i} \equiv \times_{j \neq i} \mathcal{P}^j$ is a configuration of all the players except for i , $\mathbf{p}^{-i} \in \times_{j \neq i} \Delta(\mathcal{P}^j)$ is the mixed strategies profile of all players except for i , $\varpi^{-i} \oplus a$ is the new configuration with i definitely choosing the action $a \in \mathcal{P}^i$, and $P(\mathbf{p}^{-i}, \varpi^{-i}) \equiv \prod_{j \neq i} p_{\varpi^j}^j$ is the occurrence probability of ϖ^{-i} according to \mathbf{p}^{-i} .

Remark: We abuse notation a little bit and consider the player costs $\lambda_{\varpi^i}^i$ as functions whose exact definition depends on the other players' strategies: In the general case of a mixed strategies profile \mathbf{p} , equation (2) is valid and expresses the expected cost of player i wrt \mathbf{p} , conditioned on the event that i chooses path ϖ^i . If the other players adopt a pure strategies profile ϖ^{-i} , we get the special form of equation (1) that expresses the exact cost of player i choosing action ϖ^i .

Remark: Concerning the players' private cost functions, instead of charging them for the *sum* of the expected costs of the resources that each of them chooses to use (call it the SUM-COST objective), we could also consider the maximum expected cost over all the resources in the strategy that each player adopts (call it the MAX-COST objective). This is also a valid objective, especially in scenarios dealing with bandwidth allocation in networks. Nevertheless, in the present survey we focus our interest on the SUM-COST objective, unless stated explicitly that we use some other objective.

A congestion game in which all players are indistinguishable (ie, they have the traffic demands and the same action set), is called **symmetric**. When each player's action set \mathcal{P}^i consists of sets of resources that comprise (simple) paths between a unique origin-destination pair of nodes (s_i, t_i) in (V, E) , we refer to a **(multicommodity) network congestion game**. If additionally all origin-destination pairs of the players coincide with a unique pair (s, t) we have a **single commodity network congestion game** and then all players share exactly the same action set. Observe that in general a single-commodity network congestion game is not necessarily symmetric because the players may have different demands and thus their cost functions will also differ.

2.1 Dealing with Selfish behavior.

Fix an arbitrary (mixed in general) strategies profile \mathbf{p} for a congestion game $((w_i)_{i \in N}, (\mathcal{P}^i)_{i \in N}, (d_e)_{e \in E})$. We say that \mathbf{p} is a **Nash Equilibrium (NE)** if and only if

$$\forall i \in N, \forall \alpha, \beta \in \mathcal{P}^i, p_\alpha^i > 0 \Rightarrow \lambda_\alpha^i(\mathbf{p}) \leq \lambda_\beta^i(\mathbf{p}).$$

A configuration $\varpi \in \mathcal{P}$ is a **Pure Nash Equilibrium (PNE)** if and only if

$$\forall i \in N, \forall \alpha \in \mathcal{P}^i, \lambda^i(\varpi) = \lambda_{\varpi^i}(\varpi) \leq \lambda_\alpha(\varpi^{-i} \oplus \alpha) = \lambda^i(\varpi^{-i} \oplus \alpha).$$

The **social cost** $SC(\mathbf{p})$ in this congestion game is

$$SC(\mathbf{p}) = \sum_{\varpi \in \Pi} P(\mathbf{p}, \varpi) \cdot \max_{i \in N} \{\lambda_{\varpi^i}(\varpi)\} \quad (3)$$

where $P(\mathbf{p}, \varpi) \equiv \prod_{i=1}^n p_{\varpi^i}^i$ is the probability of configuration ϖ occurring, wrt the mixed strategies profile \mathbf{p} . The **social optimum** of this game is defined as

$$OPT = \min_{\varpi \in \Pi} \left\{ \max_{i \in N} [\lambda_{\varpi^i}(\varpi)] \right\} \quad (4)$$

The **price of anarchy** for this game is then defined as

$$\mathcal{R} = \max_{\mathbf{p} \text{ is a NE}} \left\{ \frac{SC(\mathbf{p})}{OPT} \right\} \quad (5)$$

2.2 Potential Games.

Fix an arbitrary game in strategic form $\Gamma = (N, (\mathcal{P}^i)_{i \in N}, (U^i)_{i \in N})$ and some vector $\mathbf{b} \in \mathbb{R}_{>0}^n$. A function $\Phi : \mathcal{P} \rightarrow \mathbb{R}$ is called:

- an **ordinal potential** for Γ , if $\forall \varpi \in \mathcal{P}, \forall i \in N, \forall \alpha \in \mathcal{P}^i, \text{sign} [\lambda^i(\varpi) - \lambda^i(\varpi^{-i} \oplus \alpha)] = \text{sign} [\Phi(\varpi) - \Phi(\varpi^{-i} \oplus \alpha)]$,
- a **b-potential** for Γ , if $\forall \varpi \in \mathcal{P}, \forall i \in N, \forall \alpha \in \mathcal{P}^i, \lambda^i(\varpi) - \lambda^i(\varpi^{-i} \oplus \alpha) = b_i \cdot [\Phi(\varpi) - \Phi(\varpi^{-i} \oplus \alpha)]$, or
- an **exact potential** for Γ , if it is a **1**-potential for Γ .




Figure 1: An example of a layered network.

2.3 Configuration Paths and Discrete Dynamics Graph.

For a congestion game $\Gamma = ((w_i)_{i \in N}, (\mathcal{P}^i)_{i \in N}, (d_e)_{e \in E})$, a **path** in \mathcal{P} is a sequence of configurations $\gamma = (\varpi(0), \varpi(1), \dots, \varpi(k))$ s.t. $\forall j \in [k]$, $\varpi(j) = (\varpi(j-1))^{-i} \oplus \pi_i$, for some $i \in N$ and $\pi_i \in \mathcal{P}^i$. γ is a **closed path** if $\varpi(0) = \varpi(k)$. It is a **simple path** if no configuration is contained in it more than once. γ is an **improvement path** wrt Γ , if $\forall j \in [k]$, $\lambda^{i_j}(\varpi(j)) < \lambda^{i_j}(\varpi(j-1))$ where i_j is the unique player differing in its strategy between $\varpi(j)$ and $\varpi(j-1)$. That is, the unique defector of the j^{th} move in γ is actually willing to make this move because it improves its own cost. The **Nash Dynamics Graph** of Γ is a directed graph whose vertices are configurations and there is an arc from a configuration ϖ to a configuration $\varpi^{-i} \oplus \alpha$ for some $\alpha \in \mathcal{P}^i$ if and only if $\lambda^i(\varpi) > \lambda^i(\varpi^{-i} \oplus \alpha)$. The set of best replies of a player i against a configuration $\varpi^{-i} \in \mathcal{P}^{-i}$ is defined as $BR_i(\varpi^{-i}) = \arg \max_{\alpha \in \mathcal{P}^i} \{\lambda^i(\varpi^{-i} \oplus \alpha)\}$. Similarly, the set of best replies against a mixed profile \mathbf{p}^{-i} is $BR_i(\mathbf{p}^{-i}) = \arg \max_{\alpha \in \mathcal{P}^i} \{\lambda_\alpha^i(\mathbf{p}^{-i} \oplus \alpha)\}$. A path γ is a **best-reply improvement path** if each defector jumps to a best-reply pure strategy. The **Best Response Dynamics Graph** is a directed graph whose vertices are configurations and there is an arc from a configuration ϖ to a configuration $\varpi^{-i} \oplus \alpha$ for some $\alpha \in \mathcal{P}^i \setminus \{\varpi^i\}$ if and only if $\alpha \in BR_i(\varpi^{-i})$ and $\varpi^i \notin BR_i(\varpi^{-i})$.

A (finite) strategic game Γ possesses the **Finite Improvement Property (FIP)** if any improvement path of Γ has finite length. Γ possesses the **Finite Best Reply Property (FBRP)** if every best-reply improvement path is of finite length.

2.4 Isomorphism of Strategic Games.

Two games in strategic form $\Gamma = (N, (\mathcal{P}^i)_{i \in N}, (U^i)_{i \in N})$ and $\tilde{\Gamma} = (N, (\tilde{\mathcal{P}}^i)_{i \in N}, (\tilde{U}^i)_{i \in N})$ are called **isomorphic** if there exist bijections $g : \times_{i \in N} \mathcal{P}^i \mapsto \times_{i \in N} \tilde{\mathcal{P}}^i$ and $\tilde{g} : \times_{i \in N} \tilde{\mathcal{P}}^i \mapsto \times_{i \in N} \mathcal{P}^i$ s.t. $\forall \varpi \in \times_{i \in N} \mathcal{P}^i, \forall i \in N, U^i(\varpi) = \tilde{U}^i(g(\varpi))$ and $\forall \tilde{\varpi} \in \times_{i \in N} \tilde{\mathcal{P}}^i, \forall i \in N, \tilde{U}^i(\tilde{\varpi}) = U^i(\tilde{g}(\tilde{\varpi}))$.

2.5 Layered Networks.

We consider a special family of networks whose behavior wrt the price of anarchy, as we shall see, is asymptotically equivalent to that of the parallel links model of [23] (which is actually a 1-layered network): Let $\ell \geq 1$ be an integer. A directed network $G = (V, E)$ with

a distinguished source - destination pair (s, t) , $s, t \in V$, is an **ℓ -layered network** if every (simple) directed $s - t$ path has length exactly ℓ and each node lies on a directed $s - t$ path. In a layered network there are no directed cycles and all directed paths are simple. In the following, we always use $m = |E|$ to denote the number of edges in an ℓ -layered network $G = (V, E)$.

3 Related Work

3.1 Existence and tractability of PNE.

It is already known that the class of unweighted (atomic) congestion games (ie, players have the same demands and thus, the same affection on the resource delay functions) is guaranteed to have at least one PNE: actually, Rosenthal ([31]) proved that any potential game has at least one PNE and it is easy to write any unweighted congestion game as an exact potential game using Rosenthal's potential function² (eg, [12, Thm1]). In [12] it is proved that a PNE for any unweighted single-commodity network congestion game³ (no matter what resource delay functions are considered, so long as they are non-decreasing with loads) can be constructed in polynomial time, by computing the optimum of Rosenthal's potential function, through a nice reduction to min-cost flow. On the other hand, it is shown that even for a symmetric congestion game or an unweighted multicommodity network congestion game, it is PLS-complete to find a PNE (though it certainly exists).

The special case of single-commodity, parallel-edges network congestion game where the resources are considered to behave as parallel machines, has been extensively studied in recent literature. In [14] it was shown that for the case of players with varying demands and uniformly related parallel machines, there is always a PNE which can be constructed in polynomial time. It was also shown that it is NP-hard to construct the best or the worst PNE. In [17] it was proved that the fully mixed NE (FMNE), introduced and thoroughly studied in [27], is worse than any PNE, and any NE is at most $(6 + \varepsilon)$ times worse than the FMNE, for varying players and identical parallel machines. In [26] it was shown that the FMNE is the worst possible for the case of two related machines and tasks of the same size. In [25] it was proved that the FMNE is the worst possible when the global objective is the sum of squares of loads.

[13] studies the problem of constructing a PNE from any initial configuration, of social cost at most equal to that of the initial configuration. This immediately implies the existence of a PTAS for computing a PNE of minimum social cost: first compute a configuration of social cost at most $(1 + \varepsilon)$ times the social optimum ([18]), and consequently transform it into a PNE of at most the same social cost. In [11] it is also shown that even for the unrelated parallel machines case a PNE always exists, and a potential-based argument proves a convergence time (in case of integer demands) from arbitrary initial configuration to a PNE in time $\mathcal{O}(mW_{\text{tot}} + 4^{W_{\text{tot}}/m+w_{\text{max}}})$ where $W_{\text{tot}} = \sum_{i \in N} w_i$ and $w_{\text{max}} = \max_{i \in N} \{w_i\}$.

[28] studies the problem of weighted parallel-edges network congestion games with player-specific costs: each allowable action of a player consists of a single resource and each player has its own private cost function for each resource. It is shown that: (1) weighted (parallel-edges network) congestion games involving only two players, or only two possible actions for all the

²For more details on Potential Games, see [29].

³Since [12] only considers unit-demand players, this is also a symmetric network congestion game.

players, or equal delay functions (and thus, equal weights), always possess a PNE; (2) even a single-commodity, 3-player, 3-actions, weighted (parallel-edges network) congestion game may not possess a PNE (using 3-wise linear delay functions).

3.2 Price of Anarchy in Congestion Games.

In the seminal paper [23] the notion of coordination ratio, or price of anarchy, was introduced as a means for measuring the performance degradation due to lack of players' coordination when sharing common resources. In this work it was proved that the price of anarchy is $3/2$ for two related parallel machines, while for m machines and players of varying demands, $\mathcal{R} = \Omega\left(\frac{\log m}{\log \log m}\right)$ and $\mathcal{R} = \mathcal{O}(\sqrt{m \log m})$. For m identical parallel machines, [27] proved that $\mathcal{R} = \Theta\left(\frac{\log m}{\log \log m}\right)$ for the FMNE, while for the case of m identical parallel machines and players of varying demands it was shown in [22] that $\mathcal{R} = \Theta\left(\frac{\log m}{\log \log m}\right)$. In [9] it was finally shown that $\mathcal{R} = \Theta\left(\frac{\log m}{\log \log \log m}\right)$ for the general case of related machines and players of varying demands. [8] presents a thorough study of the case of general, monotone delay functions on parallel machines, with emphasis on delay functions from queuing theory. Unlike the case of linear cost functions, they show that the price of anarchy for non-linear delay functions in general is far worse and often even unbounded.

In [32] the price of anarchy in a multicommodity network congestion game among infinitely many players, each of negligible demand, is studied. The social cost in this case is expressed by the total delay paid by the whole flow in the system. For linear resource delays, the price of anarchy is at most $4/3$. For general, continuous, non-decreasing resource delay functions, the total delay of any Nash flow is at most equal to the total delay of an optimal flow for double flow demands. [33] proves that for this setting, it is actually the class of allowable latency functions and not the specific topology of a network that determines the price of anarchy.

4 Unweighted Congestion Games: [31, 29]

In this section we present some fundamental results connecting the classes of unweighted congestion games and (exact) potential games [31]. Since we refer to players of identical (say, unit) weights, the players' cost functions are

$$\lambda^i(\varpi) \equiv \sum_{e \in \varpi^i} d_e(x_e(\varpi)),$$

where, $x_e(\varpi)$ indicates the *number* of players willing to use resource e wrt configuration $\varpi \in \mathcal{P}$. The following theorem proves the strong connection of congestion games with the exact potential games.

Theorem 1 ([31, 29]) *Every (unweighted) congestion game is an exact potential game.*

Proof: Fix an arbitrary (unweighted) congestion game $\Gamma = (N, E, (\mathcal{P}^i)_{i \in N}, (d_e)_{e \in E})$. For any configuration $\varpi \in \mathcal{P}$, consider the function $\Phi(\varpi) = \sum_{e \in \bigcup_{i \in N} \varpi^i} \sum_{k=1}^{x_e(\varpi)} d_e(k)$, which we shall call **Rosenthal's potential**. We can easily show that Φ is an *exact potential* for Γ : For

this, consider arbitrary configuration $\varpi \in \mathcal{P}$, an arbitrary player $i \in N$ and an alternative (pure) strategy $\alpha \in \mathcal{P}^i \setminus \{\varpi^i\}$ for this player. Let also $\hat{\varpi} = \varpi^{-i} \oplus \alpha$. Then,

$$\begin{aligned}
\Phi(\hat{\varpi}) - \Phi(\varpi) &= \sum_{e \in \cup_j \hat{\varpi}^j} \sum_{k=1}^{x_e(\hat{\varpi})} d_e(k) - \sum_{e \in \cup_j \varpi^j} \sum_{k=1}^{x_e(\varpi)} d_e(k) \\
&= \sum_{e \in \cup_i \hat{\varpi}^i \setminus \varpi^i} \left[\sum_{k=1}^{x_e(\varpi)+1} d_e(k) - \sum_{k=1}^{x_e(\varpi)} d_e(k) \right] + \sum_{e \in \cup_i \varpi^i \setminus \hat{\varpi}^i} \left[\sum_{k=1}^{x_e(\varpi)-1} d_e(k) - \sum_{k=1}^{x_e(\varpi)} d_e(k) \right] \\
&= \sum_{e \in \hat{\varpi}^i \setminus \varpi^i} d_e(x_e(\varpi) + 1) - \sum_{e \in \varpi^i \setminus \hat{\varpi}^i} d_e(x_e(\varpi)) \\
&= \sum_{e \in \hat{\varpi}^i} d_e(x_e(\hat{\varpi})) - \sum_{e \in \varpi^i} d_e(x_e(\varpi)) = \lambda^i(\hat{\varpi}) - \lambda^i(\varpi)
\end{aligned}$$

where, we have exploited the fact that $\forall e \in E \setminus (\varpi^i \cup \hat{\varpi}^i)$ and $\forall e \in \varpi^i \cap \hat{\varpi}^i$ the load of each of these resources (ie, the number of players using them) remains the same. Additionally, $\forall e \in \hat{\varpi}^i \setminus \varpi^i$, $x_e(\hat{\varpi}) = x_e(\varpi) + 1$ and $\forall e \in \varpi^i \setminus \hat{\varpi}^i$, $x_e(\hat{\varpi}) = x_e(\varpi) - 1$. \blacksquare

Remark: The existence of a (not necessarily exact) potential for any game in strategic form directly implies the existence of a PNE for this game. The existence of an exact potential may help (as we shall see later) the *efficient* construction of a PNE, but this is not true in general.

More interestingly, Monderer and Shapley [29] proved that every (finite) potential game is isomorphic to an unweighted congestion game. The proof presented here is a new one provided by the authors of the survey. The main idea is based on the proof of Monderer and Shapley, yet it is much simpler and easier to follow.

Theorem 2 ([29]) *Every finite (exact) potential game is isomorphic to an unweighted congestion game.*

Proof: Consider an arbitrary (finite) strategic game $\Gamma = (N, (Y^i)_{i \in N}, (U^i)_{i \in N})$ among $n = |N|$ players, that admits an exact potential $\Phi : \times_{i \in N} Y^i \mapsto \mathbb{R}$. Suppose that $\forall i \in N, Y^i = \{1, 2, \dots, m_i\} \equiv [m_i]$ for some finite integer $m_i \geq 2$ (players having a unique allowable action can be safely removed from the game) and let $Y \equiv \times_{i \in N} Y^i$ be the actions space of the game. We want to construct a proper unweighted congestion game $C = (N, E, (\mathcal{P}^i)_{i \in N}, (d_e)_{e \in E})$ and a bijection $\varpi : Y \mapsto \mathcal{P}$ (recall that $\mathcal{P} \equiv \times_{i \in N} \mathcal{P}^i$ is the action space of C) such that

$$\forall i \in N, \forall \mathbf{s} \in Y, \lambda^i(\varpi(\mathbf{s})) = U^i(\mathbf{s}). \quad (6)$$

First of all, observe that for any (unweighted) congestion game C , we can express the players' costs as follows:

$$\begin{aligned}
\forall i \in N, \forall \varpi \in \mathcal{P}, \lambda^i(\varpi) &= \sum_{e \in \varpi^i} d_e(x_e(\varpi)) \\
&= \sum_{e \in \varpi^i \cap (E \setminus \cup_{j \neq i} \varpi^j)} d_e(1) + \sum_{e \in \cup_{k \neq i} \varpi^i \cap \varpi^k \cap (E \setminus \cup_{j \neq i, k} \varpi^j)} d_e(2) + (7) \\
&+ \dots + \sum_{e \in \cap_{k \in N} \varpi^k} d_e(n)
\end{aligned}$$




Figure 2: The set of resources representing the binary encodings of the configurations in the potential game.

We proceed by determining the set of shared resources $E = E_1 \cup E_2$ for our congestion game. We construct two sets of resources: The first set E_1 contains resources that determine exactly what the players in the potential game actually do. The second set E_2 contains resources that represent what all other players (except for the one considered by the resource) should not do. More specifically, fix an arbitrary configuration $\mathbf{s} \in Y$ of the players wrt the game Γ . We construct the following vector $\mathbf{e}(\mathbf{s}) \in E_1$ which is nothing more than the binary representation of this configuration (see figure 2):

$$\forall k \in N, \forall j \in [m_k], e_j^k(\mathbf{s}) = \begin{cases} 1 & \text{if } s^k = j \\ 0, & \text{otherwise} \end{cases}$$

Fix now an arbitrary player $i \in N$. We define n more resources whose binary vectors represent what the other players (ie, other than player i) in the potential game *should not* do (see figure 3):

$$\forall k \in N, \forall j \in [m_k], e_j^k(\mathbf{s}, i) = \begin{cases} 0 & \text{if } k \neq i \text{ and } s^k = j \\ 1, & \text{otherwise} \end{cases}$$

It is easy to see that $\mathbf{e}(\mathbf{s}, i)$ actually says that none of the players $k \neq i$ agrees with the profile $\mathbf{s}^{-i} \in Y^{-i}$. So, we force this binary vector $\mathbf{e}(\mathbf{s}, i)$ to have definitely 0s exactly at the positions where \mathbf{s}^{-i} has 1s for any of the players except for player i , and we place 1s anywhere else. For $m = m_1 + m_2 + \dots + m_n$, let $E_1 = \{\mathbf{e}(\mathbf{s}) \in \{0, 1\}^m : \mathbf{s} \in Y\}$ and $E_2 = \{\mathbf{e}(\mathbf{s}, i) \in \{0, 1\}^m : (i \in N) \wedge (\mathbf{s}^{-i} \in Y^{-i})\}$. The set of resources in our congestion game is then $E = E_1 \cup E_2$.

We now proceed with the definition of the action sets of the players wrt C . Indeed, $\forall i \in N$, $\mathcal{P}^i = \{\pi_1^i, \dots, \pi_{m_i}^i\}$, where $\forall j \in [m_i]$, $\pi_j^i = \{\mathbf{e} \in E : e_j^i = 1\}$. Now the bijective map that we assume is almost straightforward: $\forall \mathbf{s} = (s^i)_{i \in N} \in Y$, $\varpi(\mathbf{s}) = (\pi_{s^i}^i)_{i \in N}$.

Some crucial observations are the following: First, $\forall \mathbf{s} \in Y$, the resource $\mathbf{e}(\mathbf{s})$ is the only resource in E_1 that is used by *all* the players when the configuration $\varpi(\mathbf{s})$ is considered. For any other configuration $\mathbf{z} \in Y \setminus \{\mathbf{s}\}$, the resource $\mathbf{e}(\mathbf{s})$ is *not* used by at least one player, assuming $\varpi(\mathbf{z})$. Similarly, $\forall \mathbf{s} \in Y$, $\forall i \in N$, the resource $\mathbf{e}(\mathbf{s}, i)$ is the only resource in E_2 that is *exclusively used* by player i , assuming the configuration $\varpi(\mathbf{s})$.




Figure 3: The set of resources representing the binary encodings of the forbidden configuration for the other players in the potential game.

Now it is very simple to set the resource delay functions d_e in such a way that we assure equality (6). Indeed, observe that $\forall i \in N, \forall \mathbf{s} \in Y, \forall \alpha \in Y^i \setminus \{s^i\}$,

$$\begin{aligned} U^i(\mathbf{s}) - U^i(\mathbf{s}^{-i} \oplus \alpha) &= \Phi(\mathbf{s}) - \Phi(\mathbf{s}^{-i} \oplus \alpha) \Leftrightarrow \\ Q^i(\mathbf{s}^{-i}) \equiv U^i(\mathbf{s}) - \Phi(\mathbf{s}) &= U^i(\mathbf{s}^{-i} \oplus \alpha) - \Phi(\mathbf{s}^{-i} \oplus \alpha). \end{aligned}$$

That is, the quantity $Q^i(\mathbf{s}^{-i})$ is invariant of player i 's strategy. Recall the form of player i 's cost given in equality (7). The charging scheme of the resources that we adopt is as follows:

$$\forall \mathbf{e} \in E, d_e(k) = \begin{cases} Q^i(\mathbf{s}^{-i}), & \text{if } (k = 1) \wedge (\mathbf{e} = \mathbf{e}(\mathbf{s}, i) \in E_2) \\ 0, & \text{if } (1 < k < n) \vee (k = 1 \wedge \mathbf{e} \in E_1) \\ \Phi(\mathbf{s}), & \text{if } (k = n) \wedge (\mathbf{e} = \mathbf{e}(\mathbf{s}) \in E_1) \end{cases}$$

Now it is easy to see that for any player $i \in N$ and any configuration $\mathbf{s} \in Y$, the only resource used exclusively by i wrt to \mathbf{s} which has non-zero delay, is $\mathbf{e}(\mathbf{s}, i)$. Similarly, the only resource used by all the players that has non-zero delay is $\mathbf{e}(\mathbf{s})$. All other resources charge zero delays, no matter how many players use them. Thus,

$$\forall i \in N, \forall \mathbf{s} \in Y, (7) \Rightarrow \lambda^i(\varpi(\mathbf{s})) = d_{\mathbf{e}(\mathbf{s}, i)}(1) + d_{\mathbf{e}(\mathbf{s})}(n) = Q^i(\mathbf{s}^{-i}) + \Phi(\mathbf{s}) = U^i(\mathbf{s}).$$

■

Remark: The size of the congestion game that we use to represent a potential game is at most $(|N| + 1)$ times larger than the size of the potential game. Since an unweighted congestion game is itself an exact potential game, this implies an essential equivalence of exact potential and unweighted congestion games.

5 Existence and Complexity of Constructing PNE

In the present section we deal with issues concerning the existence and complexity of constructing PNE in weighted congestion games. Our main references for this section are

[12, 24, 15]. We start with some complexity issues concerning the construction of PNE in unweighted congestion games (in which a PNE always exists) and consequently we study existence and complexity issues in weighted congestion games in general.

Fix an arbitrary (weighted in general) congestion game $(N, E, (\mathcal{P}^i)_{i \in N}, (w_i)_{i \in N}, (d_e)_{e \in E})$ where the w_i 's denote the (positive) weights of the players.

A crucial class of problems containing the weighted network congestion games is PLS [21] (stands for *Polynomial Local Search*). This is the subclass of total functions in NP that are guaranteed to have a solution because of the fact that “every finite directed acyclic graph has a sink.” A problem Π in PLS is given by

- (a) a set of instances $I = \Sigma^*$;
- (b) $\forall x \in I$, a set of *feasible solutions* $F_x \subseteq \Sigma^{\text{poly}(|x|)}$;
- (c) a *polynomial oracle* c which, given $x \in I$, $s \in \Sigma^{\text{poly}(|x|)}$ it determines whether $s \in F_x$ and if so, it computes an integer function $c(x, s)$ (considered to be the *payoff* of s);
- (d) $\forall x \in I, s \in F_x$, a *neighborhood* of s in F_x , $N_x(s) \subseteq F_x$, and a polynomial function g which either returns a solution $s' \in N_x(s)$ s.t. $c(x, s')$ is *preferable* to $c(x, s)$ (eg, $c(x, s') < c(x, s)$ for a minimization problem), or returns **no** if no such solution exists in the neighborhood of s in F_x .

An instance of Π is then as follows: Given any $x \in I$, find a local optimum in F_x , ie, a solution $s \in F_x$ s.t. $g(s) = \mathbf{no}$.

The problem of constructing a PNE for a weighted congestion game is in PLS, in the following cases: (1) for any unweighted congestion game, since it is an exact potential game (see theorem 1), and (2) for any weighted network congestion game with linear resource delays, which admits (as we shall prove in theorem 6) a **b**-potential with $b_i = \frac{1}{2w_i}, \forall i \in N$, and thus finding PNE is equivalent to finding local optima for the optimization problem with state space the action space of the game and objective the potential of the game. On the other hand, this does not imply a polynomial-time algorithm for constructing a PNE, since (as we shall see more clearly in the weighted case) the improvements in the potential can be very small and too many. Additionally, although problems in PLS admit a PTAS, this does not imply also a PTAS for finding ε -approximate PNE (approximation of the potential does not imply also approximation of each player's payoff).

5.1 Efficient Construction of PNE in Unweighted Congestion Games

In this subsection we shall prove that for unweighted single commodity network congestion games a PNE can be constructed in polynomial time. On the other hand, even for multicommodity network congestion games it is PLS complete to construct a PNE. The main source of this subsection is the work of Fabrikant, Papadimitriou and Talwar [12].

Theorem 3 ([12]) *There is a polynomial time algorithm for finding a PNE in any unweighted single-commodity network congestion game.*

Proof: Fix an arbitrary unweighted, single-commodity network congestion game $\Gamma = (N, E, \mathcal{P}_{s-t}, (d_e)_{e \in E})$ and let $G = (V, E)$ be the underlying network. Recall that this game admits Rosenthal's exact potential $\Phi(\varpi) = \sum_{e \in E} \sum_{k=1}^{x_e(\varpi)} d_e(k)$. The algorithm computes the




Figure 4: The reduction to a min-cost flow problem [12].

optimum value of $\Phi : \mathcal{P} \mapsto \mathbb{R}$ in the action space of the game. The corresponding configuration is thus a PNE.

The algorithm exploits a reduction to a min-cost flow problem. This reduction is done as follows: We construct a new network $G'(V', E')$ with the same set of vertices $V' = V$, while the set of arcs (ie, resources) is defined as follows (see also figure 4): $\forall e = (u, v) \in E$, we add n parallel arcs from u to v , $e'(1), \dots, e'(n) \in E'$, where these arcs have a *capacity* of 1 (ie, they allow at most one player to use each of them) and fixed delays (when used) $\forall k \in [d_e(n)], d'_{e'(k)} = d_e(k)$. Let $\mathbf{f}^* \in \mathbb{R}_{\geq 0}^{|E|}$ be a minimizer of the min-cost flow problem

$$\begin{aligned}
& \min \quad \sum_{e' \in E'} d'_{e'} \cdot f_{e'} \\
& s.t. \quad \sum_{e'=(s,u) \in E'} f_{e'} - \sum_{e'=(u,s) \in E'} f_{e'} = n \\
& \quad \sum_{e'=(u,t) \in E'} f_{e'} - \sum_{e'=(t,u) \in E'} f_{e'} = n \\
& \quad \forall v \in V \setminus \{s, t\}, \quad \sum_{e'=(v,u) \in E'} f_{e'} = \sum_{e'=(u,v) \in E'} f_{e'} \\
& \quad \forall e' \in E', \quad 0 \leq f_{e'} \leq u_{e'}
\end{aligned}$$

This problem can be solved in polynomial time [1] (eg, in time $\mathcal{O}(m \log k(m + k \log k))$ using scaling algorithms) where $m = |E|$ and $k = |V|$ are the number of arcs (ie, resources) and the number of nodes in the underlying network respectively.

It is straightforward to see that any min-cost flow in G' is integral and it corresponds to a configuration in Γ that minimizes the potential of the game (which is then a PNE). ■

On the other hand, the following theorem proves that it is not that easy to construct a PNE, even in an unweighted multicommodity network congestion game. We give this theorem with a sketch of its proof:

Theorem 4 ([12]) *It is PLS-complete to find a PNE in unweighted congestion games of the following types: (i) General congestion games. (ii) Symmetric congestion games. (iii) Multicommodity network congestion games.*

Proof: We only give a sketch of the first two cases. The PLS completeness proof of case (iii) is rather complicated and therefore is not presented in this survey. The interested reader may find it in [12].

We first prove the PLS completeness of general congestion games and consequently we show that an arbitrary congestion game can be transformed into an equivalent symmetric

congestion game. In order to prove the PLS completeness of a congestion game, we shall use the following problems:

NOTALLEQUAL3SAT: Given a set N of $\{0, 1\}$ -variables and a collection C of clauses s.t. $\forall c \in C, |c| \leq 3$, is there an assignment of values to the variables so that no clause has all its literals assigned the same value?

POSNAE3FLIP: Given an instance (N, C) of NOTALLEQUAL3SAT with positive literals only and a weight function $w : C \mapsto \mathbb{R}$, find an assignment for the variables of N , s.t. the total weight of the unsatisfied clauses and the totally satisfied (ie, with all their literals set to 1) clauses cannot be decreased by a unilateral flip of the value of any variable.

It is known that POSNAE3FLIP is PLS complete [34]. Given an instance of POSNAE3FLIP, we shall construct a congestion game $\Gamma = (N, E, (\mathcal{P}^v)_{v \in N}, (d_e)_{e \in E})$ as follows: The player set of the game is exactly the set of variables N . $\forall c \in C$, we construct two resources $e_c, e'_c \in E$ whose delay functions are $d_e(k) = w(c) \cdot \mathbb{I}_{\{k=3\}}$ and $d_{e'}(k) = w(c) \cdot \mathbb{I}_{\{k=3\}}$. That is, resource e_c (or e'_c) has delay $w(c)$ only when all the 3 players it contains actually use it. Each player $v \in N$ has exactly two allowable actions indicating the possible **true** or **false** values of the corresponding variable: $\mathcal{P}^v = \{\{e_c \in E : v \in c\}, \{e'_c \in E : v \in c\}\}$. Smaller clauses (ie, of two literals) are implemented similarly. Clearly, a flip of a variable corresponds to the change in the strategy of the corresponding player. The changes in the total weight due to a flip equal the changes in the cumulative delay over all the resources. Thus, any PNE of the congestion game Γ is a local optimum (ie, a solution) of the POSNAE3FLIP problem, and vice versa.

We now proceed to show that any unweighted congestion game can be transformed into a *symmetric* congestion game. Fix again an arbitrary congestion game $\Gamma = (N, E, (\mathcal{P}^v)_{v \in N}, (d_e)_{e \in E})$. We construct an equivalent symmetric congestion game $\hat{\Gamma} = (N, \hat{E}, \hat{\mathcal{P}}, (\hat{d}_e)_{e \in \hat{E}})$ as follows: First of all we add to the set of resources n new distinct resources: $\hat{E} = E \cup \{e_i\}_{i \in N}$. The delays of these resources are: $\forall i \in N, \hat{d}_{e_i}(k) = M \cdot \mathbb{I}_{\{k \geq 2\}}$ for some sufficiently large constant M . The old resources maintain the same delay functions: $\forall e \in E, \forall k \geq 0, \hat{d}_e(k) = d_e(k)$. Now each player has the same action set $\hat{\mathcal{P}} = \times_{i \in N} \{a \cup \{e_i\} : a \in \mathcal{P}^i\}$. Observe that (by setting the constant M sufficiently large) in any PNE of $\hat{\Gamma}$ each of the distinct resources $\{e_i\}_{i \in N}$ is used by exactly one player (these resources act as if they have capacity 1 and there are only n of them). So, for any PNE in $\hat{\Gamma}$ we can easily get a PNE in Γ by simply dropping the unique new resource used by each of the players. This is done by identifying the “anonymous” players of $\hat{\Gamma}$ according to the unique resource they use, and match them with the corresponding players of Γ . ■

5.2 Existence and construction of PNE in Weighted Congestion Games

In this section we deal with the existence and tractability of PNE in weighted network congestion games. First we show that it is not always the case that a PNE exists, even for a weighted single-commodity network congestion game with only linear and 2-wise linear (ie, the maximum of two linear functions) resource delays. Recall that, as discussed previously, any unweighted congestion game has a PNE, for any kind of non-decreasing delays, due to the existence of an exact potential for these games. This result was independently proved by [15] and [24], based on similar constructions. In this survey we present the version of [15] due to its clarity and simplicity.




Figure 5: A weighted single-commodity network congestion game that may have no PNE. Consider two players with demands $w_1 = 1$ and $w_2 = 2$. The notation $a/b/c$ means that a load of 1 has delay a , a load of 2 has delay b and a load of 3 has delay c .

Lemma 5.1 ([15]) *There exist instances of weighted single-commodity network congestion games with resource delays being either linear or 2-wise linear functions of the loads, for which there is no PNE.*

Proof: We demonstrate this by the example shown in figure 5. In this example there are exactly two players of demands $w_1 = 1$ and $w_2 = 2$, from node s to node t . The possible paths that the two players may follow are labeled in the figure. The resource delay functions are indicated by the 3 possible values they may take given the two players. Observe now that this example has no PNE: there is a simple closed path $\gamma = \langle (P3, P2), (P3, P4), (P1, P4), (P1, P2), (P3, P2) \rangle$ of length 4 that is an improvement path (actually, each defecting player moves to its new best choice, so this is a best-reply improvement path) and additionally, any other configuration not belonging to γ is either one, or two best-reply moves away from some of these nodes. Therefore there is no sink in the Nash Dynamics Graph of the game and thus there exists no PNE. Observe that the delay functions are *not* player-specific in our example, as was the case in [28]. ■

Consequently we show that there may exist no exact potential function for a weighted single-commodity network congestion game, even when the resource delays are identical to their loads. The next argument shows that theorem 1 does not hold anymore even in this simplest case of weighted congestion games.

Lemma 5.2 ([15]) *There exist weighted single-commodity network congestion games which are not exact potential games, even when the resource delays are identical to their loads.*

Proof: Let $\Gamma = ((w_i)_{i \in N}, (\mathcal{P}^i)_{i \in N}, (d_e)_{e \in E})$ denote a weighted single commodity network congestion game with $d_e(x) = x$, $\forall e \in E$. Recall the definition of players' costs for a configuration (eq. (1)). Let's now define the quantity $I(\gamma, \lambda) = \sum_{k=1}^r [\lambda^{i_k}(\varpi(k)) - \lambda^{i_k}(\varpi(k-1))]$, where i_k is the unique player in which the configurations $\varpi(k)$ and $\varpi(k-1)$ differ. Our proof is based on the fact that Γ is an (exact) potential game if and only if every simple closed path γ of length 4 has $I(\gamma, \lambda) = 0$ ([29, Thm2.8]).

For the sake of contradiction, assume that every closed simple path γ of length 4 for Γ has $I(\gamma, \lambda) = 0$, fix arbitrary configuration ϖ and consider the path $\gamma = (\varpi, x = \varpi^{-1} \oplus \pi_1, y = \varpi^{-(1,2)} \oplus (\pi_1, \pi_2), z = \varpi^{-2} \oplus \pi_2, \varpi)$ for some paths $\pi_1 \neq \varpi^1$ and $\pi_2 \neq \varpi^2$. We shall

demonstrate that $I(\gamma, \lambda)$ cannot be identically 0 when there are at least two players of different demands. So, consider that the first two players have different demands: $w_1 \neq w_2$. We observe that

$$\lambda^1(x) - \lambda^1(\varpi) = \sum_{e \in \pi_1} \theta_e(x) - \sum_{e \in \varpi^1} \theta_e(\varpi) = |\pi_1 \setminus \varpi^1| \cdot w_1 + \sum_{e \in \pi_1 \setminus \varpi^1} \theta_e(\varpi) - \sum_{e \in \varpi^1 \setminus \pi_1} \theta_e(\varpi)$$

since the resources in $\varpi^1 \cap \pi_1$ retain their initial loads. Similarly we have:

$$\begin{aligned} \lambda^2(y) - \lambda^2(x) &= \sum_{e \in \pi_2 \setminus \varpi^2} [\theta_e(x) + w_2] - \sum_{e \in \varpi^2 \setminus \pi_2} \theta_e(x) \\ &= |\pi_2 \setminus \varpi^2| \cdot w_2 + \sum_{e \in \pi_2 \setminus \varpi^2} \theta_e(x) - \sum_{e \in \varpi^2 \setminus \pi_2} \theta_e(x) \\ \lambda^1(z) - \lambda^1(y) &= \sum_{e \in \varpi^1 \setminus \pi_1} \theta_e(z) - \sum_{e \in \pi_1 \setminus \varpi^1} [\theta_e(z) + w_1] \\ &= \sum_{e \in \varpi^1 \setminus \pi_1} \theta_e(z) - \sum_{e \in \pi_1 \setminus \varpi^1} \theta_e(z) - |\pi_1 \setminus \varpi^1| \cdot w_1 \\ \lambda^2(\varpi) - \lambda^2(z) &= \sum_{e \in \varpi^2 \setminus \pi_2} \theta_e(\varpi) - \sum_{e \in \pi_2 \setminus \varpi^2} \theta_e(\varpi) - |\pi_2 \setminus \varpi^2| \cdot w_2 \end{aligned}$$

Thus, since $I \equiv I(\gamma, \lambda) = \lambda^1(x) - \lambda^1(\varpi) + \lambda^2(y) - \lambda^2(x) + \lambda^1(z) - \lambda^1(y) + \lambda^2(\varpi) - \lambda^2(z)$, we get:

$$\begin{aligned} I &= \sum_{e \in \pi_1 \setminus \varpi^1} [\theta_e(\varpi) - \theta_e(z)] + \sum_{e \in \pi_2 \setminus \varpi^2} [\theta_e(x) - \theta_e(\varpi)] + \\ &+ \sum_{e \in \varpi^1 \setminus \pi_1} [\theta_e(z) - \theta_e(\varpi)] + \sum_{e \in \varpi^2 \setminus \pi_2} [\theta_e(\varpi) - \theta_e(x)] \end{aligned}$$

Observe now that

$$\begin{aligned} \forall e \in \pi_1 \setminus \varpi^1, \quad \theta_e(\varpi) - \theta_e(z) &= \theta_e(\varpi) - \theta_e(\varpi^{-2} \oplus \pi_2) = w_2 \cdot (\mathbb{I}_{\{e \in \varpi^2 \setminus \pi_2\}} - \mathbb{I}_{\{e \in \pi_2 \setminus \varpi^2\}}) \\ \forall e \in \pi_2 \setminus \varpi^2, \quad \theta_e(x) - \theta_e(\varpi) &= \theta_e(\varpi^{-1} \oplus \pi_1) - \theta_e(\varpi) = w_1 \cdot (\mathbb{I}_{\{e \in \pi_1 \setminus \varpi^1\}} - \mathbb{I}_{\{e \in \varpi^1 \setminus \pi_1\}}) \\ \forall e \in \varpi^1 \setminus \pi_1, \quad \theta_e(z) - \theta_e(\varpi) &= \theta_e(\varpi^{-2} \oplus \pi_2) - \theta_e(\varpi) = w_2 \cdot (\mathbb{I}_{\{e \in \pi_2 \setminus \varpi^2\}} - \mathbb{I}_{\{e \in \varpi^2 \setminus \pi_2\}}) \\ \forall e \in \varpi^2 \setminus \pi_2, \quad \theta_e(\varpi) - \theta_e(x) &= \theta_e(\varpi) - \theta_e(\varpi^{-1} \oplus \pi_1) = w_1 \cdot (\mathbb{I}_{\{e \in \varpi^1 \setminus \pi_1\}} - \mathbb{I}_{\{e \in \pi_1 \setminus \varpi^1\}}) \end{aligned}$$

Then, $I = (w_1 - w_2) \cdot [|(\pi_1 \setminus \varpi^1) \cap (\pi_2 \setminus \varpi^2)| + |(\varpi^1 \setminus \pi_1) \cap (\varpi^2 \setminus \pi_2)| - |(\varpi^1 \setminus \pi_1) \cap (\pi_2 \setminus \varpi^2)| - |(\pi_1 \setminus \varpi^1) \cap (\varpi^2 \setminus \pi_2)|]$, which is typically not equal to zero for a single-commodity network. It should be noted that the second parameter, which is network dependent, can be non-zero even for some cycle of a very simple network. For example, in the network of figure 5 (which is a simple 2-layered network) the simple closed path $\gamma = (\varpi(0) = (P1, P3), \varpi(1) = (P2, P3), \varpi(2) = (P2, P1), \varpi(3) = (P1, P1), \varpi(4) = (P1, P3))$ has this quantity equal to -4 and thus no weighted single commodity network congestion game on this network can admit an exact potential. \blacksquare

Our next step is to focus our interest on the ℓ -layered networks with resource delays identical to their loads. We shall prove that any weighted ℓ -layered network congestion game

with these delays admits at least one PNE, which can be computed in pseudo-polynomial time. Although we already know that even the case of weighted ℓ -layered network congestion games with delays equal to the loads cannot have any exact potential⁴, we will next show that $\Phi(\varpi) \equiv \sum_{e \in E} [\theta_e(\varpi)]^2$ is a \mathbf{b} -potential for such a game and some positive n -vector \mathbf{b} , assuring the existence of a PNE.

Theorem 5 *For any weighted ℓ -layered network congestion game with resource delays equal to their loads, at least one PNE exists and can be computed in pseudo-polynomial time.*

Proof: Fix an arbitrary ℓ -layered network (V, E) and denote by \mathcal{P} all the s – t paths in it from the unique source s to the unique destination t . Let $\varpi \in \mathcal{P}^n$ be an arbitrary configuration of the players for the corresponding congestion game on (V, E) . Also, let i be a user of demand w_i and fix some path $\alpha \in \mathcal{P}$. Denote $\hat{\varpi} \equiv \varpi^{-i} \oplus \alpha$. Observe that

$$\begin{aligned}
\Phi(\varpi) - \Phi(\hat{\varpi}) &= \sum_{e \in E} (\theta_e^2(\varpi) - \theta_e^2(\hat{\varpi})) \\
&= \sum_{e \in \varpi^i \setminus \alpha} (\theta_e^2(\varpi) - \theta_e^2(\hat{\varpi})) + \sum_{e \in \alpha \setminus \varpi^i} (\theta_e^2(\varpi) - \theta_e^2(\hat{\varpi})) \\
&= \sum_{e \in \varpi^i \setminus \alpha} ([\theta_e(\varpi^{-i}) + w_i]^2 - \theta_e^2(\varpi^{-i})) + \sum_{e \in \alpha \setminus \varpi^i} (\theta_e^2(\varpi^{-i}) - [\theta_e(\varpi^{-i}) + w_i]^2) \\
&= \sum_{e \in \varpi^i \setminus \alpha} (w_i^2 + 2w_i\theta_e(\varpi^{-i})) - \sum_{e \in \alpha \setminus \varpi^i} (w_i^2 + 2w_i\theta_e(\varpi^{-i})) \\
&= w_i^2 \cdot (|\varpi^i \setminus \alpha| - |\alpha \setminus \varpi^i|) + 2w_i \cdot \left(\sum_{e \in \varpi^i \setminus \alpha} \theta_e(\varpi^{-i}) - \sum_{e \in \alpha \setminus \varpi^i} \theta_e(\varpi^{-i}) \right) \\
&= 2w_i \cdot \left(\sum_{e \in \varpi^i \setminus \alpha} \theta_e(\varpi^{-i}) - \sum_{e \in \alpha \setminus \varpi^i} \theta_e(\varpi^{-i}) \right) = 2w_i \cdot [\lambda^i(\varpi) - \lambda^i(\hat{\varpi})],
\end{aligned}$$

since, $\forall e \in \varpi^i \cap \alpha$, $\theta_e(\varpi) = \theta_e(\hat{\varpi})$, in ℓ -layered networks $|\varpi^i \setminus \alpha| = |\alpha \setminus \varpi^i|$, $\lambda^i(\varpi) = \sum_{e \in \varpi^i} \theta_e(\varpi) = \sum_{e \in \varpi^i \setminus \alpha} \theta_e(\varpi^{-i}) + w_i|\varpi^i \setminus \alpha| + \sum_{e \in \varpi^i \cap \alpha} \theta_e(\varpi)$ and $\lambda^i(\hat{\varpi}) = \sum_{e \in \alpha} \theta_e(\hat{\varpi}) = \sum_{e \in \alpha \setminus \varpi^i} \theta_e(\varpi^{-i}) + w_i|\alpha \setminus \varpi^i| + \sum_{e \in \varpi^i \cap \alpha} \theta_e(\varpi)$. Thus, Φ is a \mathbf{b} -potential for our game, where $\mathbf{b} = (1/(2w_i))_{i \in N} > \mathbf{0}$, assuring the existence of at least one PNE.

We proceed with the construction of a PNE in pseudopolynomial time. Wlog assume now that the players have integer weights. Then each player performing any improving defection, must reduce his cost by at least 1 and thus the potential function decreases by at least $2w_{\min} \geq 2$ along each arc of the Nash Dynamics Graph of the game. Consequently, the naïve algorithm that, starting from an arbitrary initial configuration $\varpi \in \mathcal{P}^n$, follows any improvement path that leads to a sink (ie, a PNE) of the Nash Dynamics Graph, cannot contain more than $\frac{1}{2}|E|W_{\text{tot}}^2$ defections, since $\forall \varpi \in \mathcal{P}^n$, $\Phi(\varpi) \leq |E|W_{\text{tot}}^2$. ■

A recent improvement, based essentially on the same technique as above, generalizes the last result to the case of arbitrary multicommodity network congestion games with linear resource delays (we state the result here without its proof):

Theorem 6 ([16]) *For any weighted multi-commodity network congestion game with linear resource delays, at least one PNE exists and can be computed in pseudo-polynomial time.*

⁴The example at the end of the proof of lemma 5.2 involves the 2-layered network of figure 5.

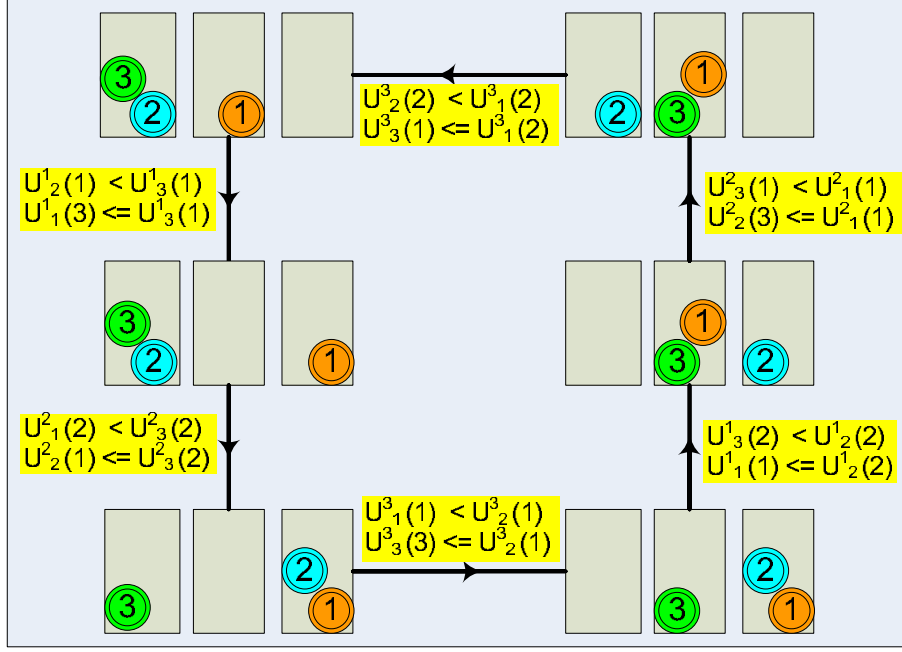


Figure 6: An example of a 3-players, 3-strategies unweighted congestion game with a best-reply cycle.

6 Congestion Games on Parallel Links with Player Specific Payoffs [28]

In his work, Milchtaich studies a variant of the classical (unweighted) congestion games, where the resource charging functions are not universal, but player-specific. On the other hand, he makes two simplifying (yet crucial) assumptions:

- (1) each player may choose only one resource from a pool E of resources (shared to all the players) for his service (ie, this is modeled as the *parallel-links* model of Koutsoupias and Papadimitriou [23]), and
- (2) the received payoff is *monotonically non-increasing* with the number of players selecting it. Although they do not always admit a potential, these games always possess a PNE.

In his paper, Milchtaich proves that unweighted congestion games on parallel links with player-specific payoffs, involving only two strategies, possess the Finite Improvement Property (FIP). It is also rather straightforward that any 2-players unweighted congestion game on parallel links with player-specific payoffs possesses the Finite Best Reply improvement path Property (FBRP).

Milchtaich also gave an example of an unweighted congestion game on 3 parallel links with 3 players, for which there is a best-reply cycle (although there is a PNE). The example is shown in figure 6. In this example, we only determine the necessary conditions on the payoff functions of the players for the existence of the best-response cycle (see figure). It is easily verified that this system of inequalities is feasible, and that configurations (3, 1, 2) and (2, 3, 1) are PNE for the game.

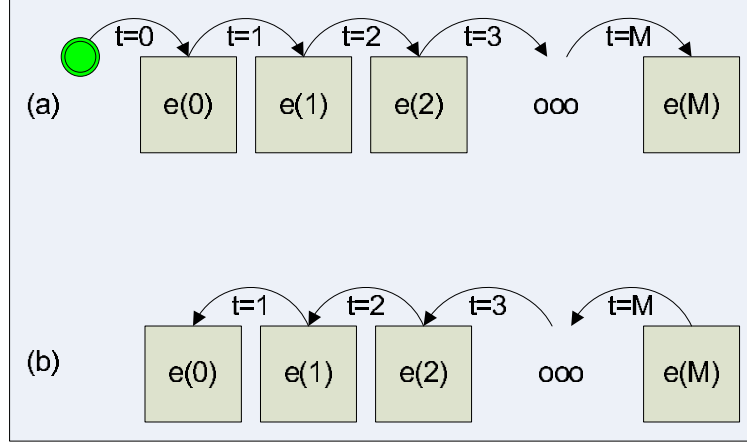


Figure 7: The two types of best-reply improvement paths considered by [28].

Theorem 7 ([28]) *Every unweighted congestion game on parallel links with player-specific, monotonically non-increasing payoffs of the resources, possesses a PNE.*

Proof: First of all, we need to prove the following lemma that bounds the lengths of best-reply paths of a specific kind.

Lemma 6.1 *Let $\langle e(0), e(1), \dots, e(M) \rangle$ be a sequence of (possibly repeated) resources from E .*

type-(a) path: *Let $\langle \varpi(1), \varpi(2), \dots, \varpi(M) \rangle$ be a best-reply improvement path of the game, s.t. $\forall t \geq 1, \forall e \in E$,*

$$x_e(t) = \begin{cases} x_e(t-1), & \text{if } e \neq e(t-1), e(t) \\ x_e(t-1) + 1, & \text{if } e = e(t) \\ x_e(t-1) - 1, & \text{if } e = e(t-1) \end{cases}$$

Then $M \leq |N|$.

type-(b) path: *Let $\langle \varpi(1), \varpi(2), \dots, \varpi(M) \rangle$ be a best-reply improvement path of the game, s.t. $\forall t \geq 1, \forall e \in E$,*

$$x_e(t) = \begin{cases} x_e(t-1), & \text{if } e \neq e(t-1), e(t) \\ x_e(t-1) + 1, & \text{if } e = e(t-1) \\ x_e(t-1) - 1, & \text{if } e = e(t) \end{cases}$$

Then, $M \leq |N| \cdot (|E| - 1)$.

Proof: (a) Let $\forall e \in E, x_e^{min} \equiv \min_{0 \leq t \leq M} x_e(t)$. By definition of the best-reply sequence, we observe that $\forall e \in E, \forall t \geq 1, x_{e(t)}(t) = x_{e(t)}(t-1) + 1 = x_e^{min} + 1$. That is, a resource reaches its maximum load exactly when it is the next node in the sequence of resources, reaches its minimum load in the very next time step, and remains at its minimum load until before it appears again in this sequence (see case (a) of figure 7). But then, the unique deviator in each move goes to a resource that reaches its maximum load, whereas all the other resources are at their minimum loads at the same time. This implies that there is no chance that the

specific player will move away from this new resource that he chose, till the end of the best reply-path under consideration. More formally, $\forall t \geq 1$, let $i(t)$ be the unique deviator that moves from $e(t-1)$ to $e(t)$. Then,

$$\forall e \in E, \forall t+1 \leq t' \leq M, U_{e(t)}^{i(t)}(x_e(t')) \geq U_{e(t)}^{i(t)}(x_{e(t)}^{\min} + 1) \geq U_e^{i(t)}(x_e^{\min}) \geq U_e^{i(t)}(x_e(t'))$$

and so, $i(t)$ cannot move away from $e(t)$ till the end of the best-reply improvement path. Since this holds for all players, this path may have length at most $|N|$.

(b) Let again $\forall e \in E, x_e^{\min} \equiv \min_{0 \leq t \leq M} x_e(t)$. We observe that $\forall 1 \leq t \leq M, x_{e(t)}(t) = x_{e(t)}^{\min}$ and $x_{e(t-1)}(t) = x_{e(t-1)}^{\min} + 1$. Due to the best-reply moves considered in the path, if $i(t)$ is again the unique deviator at time t , then $\forall 1 \leq t \leq M$,

$$U_{e(t)}^{i(t)}(x_{e(t)}(t-1)) = U_{e(t)}^{i(t)}(x_{e(t)}^{\min} + 1) < U_{e(t-1)}^{i(t)}(x_{e(t-1)}(t-1) + 1) = U_{e(t-1)}^{i(t)}(x_{e(t-1)}^{\min} + 1)$$

which implies that $\forall t' > t$ player $i(t)$ residing at a resource $e \neq e(t)$ could never prefer to deviate to $e(t)$ rather than deviate to (or, stay at, if already there) resource $e(t-1)$. That is, player $i(t)$ can never go back to the resource from which it defected once, till the end of the best-reply path. Thus, each player can make at most $|E| - 1$ moves and so, $M \leq |N| \cdot (|E| - 1)$. \blacksquare

The proof of the theorem proceeds by induction on the number of players. Trivially, for a single player we know that as soon as he is placed at its best reply, this is also a PNE (there is nothing else to move). We assume now that any unweighted congestion game on parallel links with player-specific payoffs and $|\tilde{N}| < n$, $\tilde{\Gamma} = (\tilde{N}, E, (U_e^i)_{i \in \tilde{N}, e \in E})$, possesses a PNE. We want to prove that this is also the case for any such game with n players. Let $N = [n]$ and $\Gamma = (N, E, (U_e^i)_{i \in N, e \in E})$ be such a game. We temporarily pull player n out of the game and let the remaining $n-1$ players continue playing, until they reach a PNE (without player n) $\mathbf{s} \in E^{n-1}$. That the PNE \mathbf{s} exists for the pruned game $\tilde{\Gamma} = ([n-1], E, (U_e^i)_{i \in [n-1], e \in E})$ holds by inductive hypothesis. Now we let player n be assigned to a best-reply resource $e \in BR_n(\mathbf{s})$, and we thus construct a configuration $\varpi(0)$ as follows: $\varpi^n(0) = e$; $\forall i \in [n-1], \varpi^i(0) = s^i$.

Consequently, starting from $\varpi(0)$, we construct a best-reply *maximal* improvement path of type (a) (see previous lemma) $\Pi(a) = \langle \varpi(0), \varpi(1), \dots, \varpi(M) \rangle$, which we already know that is of length at most n . We claim that the terminal configuration of this path is a PNE for Γ . Clearly, any player that deviated to a best-reply resource in $\Pi(a)$ does not move again and is also at a best-reply resource in $\varpi(M)$ (see the proof of the lemma). So we only have to consider players that have not defected during $\Pi(a)$. Fix any such player i , residing at resource $e = \varpi^i(0) = s^i$. Observe that in $\varpi(M)$ the following holds:

$$\forall e \in E, x_e(M) = \begin{cases} x_e(0), & \text{if } e \neq e(M), \\ x_e(0) + 1, & \text{if } e = e(M) \end{cases}$$

Observe that if some player $i : s^i(M) = e(M)$ that has not moved during $\Pi(a)$ is not at his best reply in $\varpi(M)$, then we can set $e(M+1) \in BR_i(\varpi^{-i}(M))$ and thus augment the best-reply improvement path $\Pi(a)$, which contradicts its maximality assumption. Consider now any player $i \in e \neq e(M)$ that has not moved during $\Pi(a)$. This player is certainly at a best-reply resource $e \in BR_i(\varpi^{-i}(M))$ since this resource has exactly the same load as in \mathbf{s} and any other resource has at least the load it had in \mathbf{s} . So, $\varpi(M)$ is a PNE for Γ since every player is at a best-reply resource wrt $\varpi(M)$. This completes the proof of the theorem. \blacksquare

Remark: Observe that the proof of this theorem is constructive, and thus also implies a path of length at most $|N|$ that leads to a PNE. But this is not necessarily an improvement path, when all players are considered to coexist all the time, and therefore there is no justification of the adoption of such a path by the (selfish) players. Milchtaich [28], using an argument of the same flavor as in the above theorem, proves that from an arbitrary initial configuration and allowing only best-reply defections, there is a best-reply improvement path of length at most $|E| \cdot \binom{|N|+1}{2}$. The idea is, starting from an arbitrary configuration $\varpi(0)$, to let the players construct a best-reply improvement path of type-(a), and then (if necessary) construct a best-reply improvement path of type-(b) starting from the terminal configuration of the previous path. It is then easily shown that the terminal configuration of the second path is a PNE of the game.

The unweighted congestion games on parallel links and with player-specific payoffs are **weakly acyclic** games, in the sense that from any initial configuration $\varpi(0)$ of the players, there is at least one best-reply improvement path connecting it to a PNE. Of course, this does not exclude the existence of best-reply cycles (see example of figure 6). But, it is easily shown that when the deviations of the players occur sequentially and in each step the next deviator is chosen randomly (among the potential deviators) to a randomly chosen best-reply resource, then this path will converge almost surely to a PNE in finite time.

6.1 Allowing different weights to the players

Milchtaich proposed a generalization of his variant of congestion games, by allowing the players to have distinct weights, denoted by a weight vector $\mathbf{w} = (w^1, w^2, \dots, w^n) \in \mathbb{R}_{>0}^n$. In that case, the (player-specific) payoff of each player on a resource $e \in E$ depends on the load $\theta_e(\varpi) \equiv \sum_{i:e \in \varpi^i} w_i$, rather than the number of players willing to use it.

For the case of weighted congestion games on parallel links with player specific payoffs, it is easy to verify (in a similar fashion as for the unweighted case) that:

- If there are only two available strategies then FIP holds.
- If there are only two players then FBRP holds.
- For the special case of resource specific payoffs, FIP holds.

On the other hand, there exists a 3-players, 3-strategies game that possesses no PNE. For example, see the instance shown in figure 8, where the three players have essentially two strategies each (a “LEFT” and a “RIGHT” strategy) to choose from, while their third strategies give them strictly minimal payoffs and can never be chosen by selfish moves. The rationale of this game is that, in principle, player 1 would like to avoid using the same link as player 3, which in turn would like to avoid using the same link as player 2, which would finally want to avoid player 1.

The inequalities shown in figure 8(b) demonstrate the necessary conditions for the existence of a best-reply cycle among 6 configurations of the players. It is easy to verify also that any other configuration has either at least one MIN strategy for some player (and in that case this player wants to defect) or is one of $(2, 2, 1)$, $(3, 3, 3)$. The only thing that remains to assure, is that $(2, 3, 1)$ is strictly better for player 2 than $(2, 2, 1)$ (ie, player 2 would like to avoid player 1) and that $(3, 3, 1)$ is better for player 3 than $(3, 3, 3)$ (ie, player 3 would like to

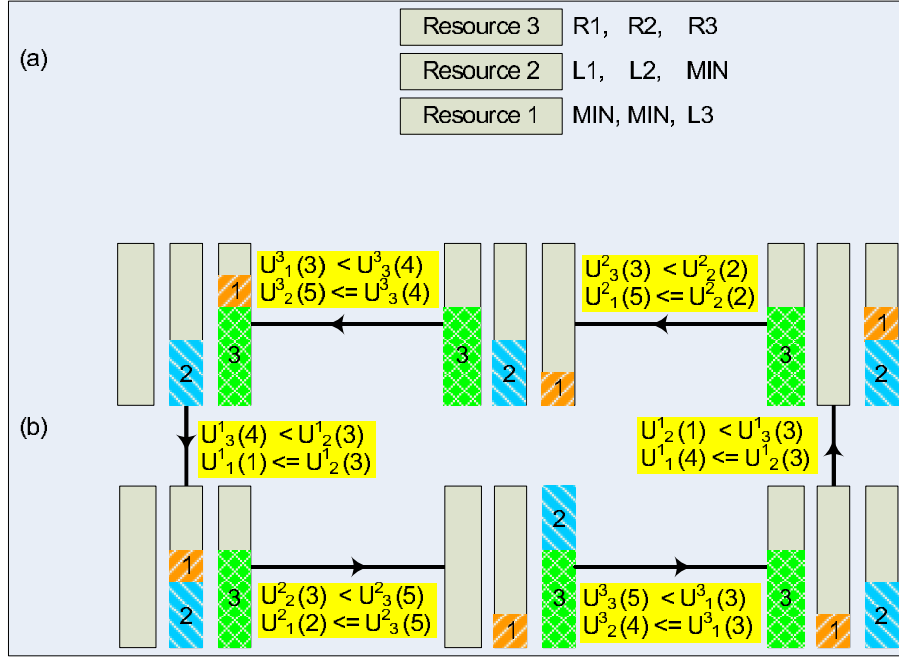


Figure 8: A 3-players weighted congestion game on 3 parallel links with player-specific payoffs without a PNE [28]. (a) The LEFT-RIGHT strategies of the players. (b) The best-reply cycle.




Figure 9: An example of a single-source network congestion game without a PNE ([32]).

avoid player 2). The feasibility of the whole system of inequalities can be trivially checked to hold, and thus this game cannot have any PNE since there is no sink in its Dynamics graph.

7 The Price of Anarchy of Weighted Congestion Games: [15]

In this section we focus our interest on weighted ℓ -layered network congestion games where the resource delays are identical to their loads. Our source for this section is [15]. This case comprises a non-trivial generalization of routing through identical parallel channels.

The main reason why we focus on this specific category of resource delays is that selfish unsplittable flows can have unbounded price of anarchy even for linear resource delays. In [32, p. 256] an example is given where the price of anarchy is unbounded (see figure 9). This example is easily converted in an ℓ -layered network. The resource delay functions used are either constant or M/M/1-like delay functions. But we can be equally bad even with linear resource delay functions: Observe the following example of figure 10. Two players, each of




Figure 10: Example of an ℓ -layered network with linear resource delays and unbounded anarchy.

unit demand, want to move selfishly from s to t . The edge delays are shown above them. We assume that $a \gg b \gg 1 \geq c$. It is easy to see that the configuration (sCBt,sADt) is a PNE of social cost $2 + b$ while the optimum configuration is (sABt,sCDt) whose social optimum is $2 + c$. Thus, $\mathcal{R} = \frac{b+2}{c+2}$.

In the following, we restrict our attention to ℓ -layered networks whose resource delays are equal to their loads. Our main tool is to interpret a strategies profile as a flow in the underlying network.

7.1 Flows and mixed strategies profiles.

Fix an arbitrary ℓ -layered network $G = (V, E)$ and n distinct players willing to satisfy their own traffic demands from the unique source $s \in V$ to the unique destination $t \in V$. Again, $\mathbf{w} = (w_i)_{i \in [n]}$ denotes the varying demands of the players. Fix an arbitrary mixed strategies profile $\mathbf{p} = (p_1, p_2, \dots, p_n)$. A *feasible flow* for the n players is a function $\rho : \mathcal{P} \mapsto \mathbb{R}_{\geq 0}$, s.t. $\sum_{\pi \in \mathcal{P}} \rho(\pi) = W_{\text{tot}} \equiv \sum_{i \in [n]} w_i$, ie, all players' demands are actually met. We distinguish between *unsplittable* and *splittable* (feasible) flows. A flow is unsplittable if each player's traffic demand is satisfied by a unique path of \mathcal{P} . A flow is splittable if the traffic demand of each player can be routed over several paths of \mathcal{P} .

We map the mixed strategies profile \mathbf{p} to a flow $\rho_{\mathbf{p}}$ as follows: For each $s - t$ path $\pi \in \mathcal{P}$, $\rho_{\mathbf{p}}(\pi) \equiv \sum_{i \in [n]} w_i \cdot p_i(\pi)$. That is, we handle the *expected load traveling along π according to \mathbf{p}* as a splittable flow, where player i routes a fraction of $p_i(\pi)$ of its total demand w_i along π . Observe that, if \mathbf{p} is actually a pure strategies profile, the corresponding flow is then unsplittable. Recall now that for each edge $e \in E$,

$$\theta_e(\mathbf{p}) \equiv \sum_{i=1}^n \sum_{\pi \ni e} w_i p_i(\pi) = \sum_{\pi \ni e} \rho_{\mathbf{p}}(\pi) \equiv \theta_e(\rho_{\mathbf{p}})$$

denotes the expected load (and in our case, also the expected delay) of e wrt \mathbf{p} . As for the expected delay along a path $\pi \in \mathcal{P}$ according to \mathbf{p} , this is

$$\theta_{\pi}(\mathbf{p}) \equiv \sum_{e \in \pi} \theta_e(\mathbf{p}) = \sum_{e \in \pi} \sum_{\pi' \ni e} \rho_{\mathbf{p}}(\pi') = \sum_{\pi' \in \mathcal{P}} |\pi \cap \pi'| \rho_{\mathbf{p}}(\pi') \equiv \theta_{\pi}(\rho_{\mathbf{p}})$$

Let $\theta^{\min}(\rho) = \min_{\pi \in \mathcal{P}} \{\theta_{\pi}(\rho)\}$ be the minimum expected delay among all $s - t$ paths. From now on for simplicity we drop the subscript of \mathbf{p} from its corresponding flow $\rho_{\mathbf{p}}$, when this is clear by the context. We evaluate flow ρ using the objective of *maximum latency*, which is defined as

$$L(\rho) \equiv \max_{\pi: \rho(\pi) > 0} \{\theta_{\pi}(\rho)\} = \max_{\pi: \exists i, p_i(\pi) > 0} \{\theta_{\pi}(\mathbf{p})\} \equiv L(\mathbf{p}) \quad (8)$$

$L(\rho)$ is nothing but the *maximum expected delay paid by the players*, wrt \mathbf{p} . From now on, we use ρ^* and ρ_f^* to denote the optimal unsplittable and splittable flows respectively.

In addition, we sometimes evaluate flow ρ using the objective of *total latency*, which is defined as

$$C(\rho) \equiv \sum_{\pi \in \mathcal{P}} \rho(\pi) \theta_\pi(\rho) = \sum_{e \in E} \theta_e^2(\rho) = \sum_{e \in E} \theta_e^2(\mathbf{p}) \equiv C(\mathbf{p}) \quad (9)$$

The second equality is obtained by summing over the edges of π and reversing the order of the summation.

7.2 Flows at Nash equilibrium.

Let \mathbf{p} be a mixed strategies profile and let ρ be the corresponding flow. For a ℓ -layered network with resource delays equal to the loads, the cost of player i on path π is $\lambda_\pi^i(\mathbf{p}) = \ell w_i + \theta_\pi^{-i}(\mathbf{p})$, where $\theta_\pi^{-i}(\mathbf{p})$ is the expected delay along path π if the demand of player i was removed from the system:

$$\theta_\pi^{-i}(\mathbf{p}) = \sum_{\pi' \in \mathcal{P}} |\pi \cap \pi'| \sum_{j \neq i} w_j p_j(\pi') = \theta_\pi(\mathbf{p}) - w_i \sum_{\pi' \in \mathcal{P}} |\pi \cap \pi'| p_i(\pi') \quad (10)$$

Thus, $\lambda_\pi^i(\mathbf{p}) = \theta_\pi(\mathbf{p}) + [\ell - \sum_{\pi' \in \mathcal{P}} |\pi \cap \pi'| p_i(\pi')] w_i$. Observe now that, if \mathbf{p} is a NE, then $L(\mathbf{p}) = L(\rho) \leq \theta^{\min}(\rho) + \ell w_{\max}$. Otherwise, the players routing their traffic on a path of expected delay greater than $\theta^{\min}(\rho) + \ell w_{\max}$ could improve their delay by defecting to a path of expected delay $\theta^{\min}(\rho)$. We sometimes say that a flow ρ corresponding to a mixed strategies profile \mathbf{p} is a NE with the understanding that it is actually \mathbf{p} which is a NE.

7.3 Maximum Latency versus Total Latency.

We show that if the resource delays are equal to their loads, a splittable flow is optimal wrt the objective of maximum latency if and only if it is optimal wrt the objective of total latency. As a corollary, we obtain that the optimal splittable flow defines a NE where all players adopt the same mixed strategy.

Lemma 7.1 *There is a unique feasible splittable flow ρ which minimizes both $L(\rho)$ and $C(\rho)$.*

Proof: For every feasible flow ρ , the average latency of ρ cannot exceed its maximum latency:

$$C(\rho) = \sum_{\pi \in \mathcal{P}} \rho(\pi) \theta_\pi(\rho) = \sum_{\pi: \rho(\pi) > 0} \rho(\pi) \theta_\pi(\rho) \leq L(\rho) W_{\text{tot}} \quad (11)$$

A splittable flow ρ minimizes $C(\rho)$ if and only if for every $\pi_1, \pi_2 \in \mathcal{P}$ with $\rho(\pi_1) > 0$, $\theta_{\pi_1}(\rho) \leq \theta_{\pi_2}(\rho)$ (e.g., [6], [30, Section 7.2], [32, Corollary 4.2]). Hence, if ρ is optimal wrt the objective of total latency, for all paths $\pi \in \mathcal{P}$, $\theta_\pi(\rho) \geq L(\rho)$. Moreover, if $\rho(\pi) > 0$, then $\theta_\pi(\rho) = L(\rho)$. Therefore, if ρ minimizes $C(\rho)$, then the average latency is equal to the maximum latency:

$$C(\rho) = \sum_{\pi \in \mathcal{P}: \rho(\pi) > 0} \rho(\pi) \theta_\pi(\rho) = L(\rho) W_{\text{tot}} \quad (12)$$

Let ρ be the feasible splittable flow that minimizes the total latency and let ρ' be the feasible splittable flow that minimizes the maximum latency. We prove the lemma by establishing that the two flows are identical.

Observe that $L(\rho') \geq \frac{C(\rho')}{W_{\text{tot}}} \geq \frac{C(\rho)}{W_{\text{tot}}} = L(\rho)$. The first inequality follows from Ineq. (11), the second from the assumption that ρ minimizes the total latency and the last equality from Eq. (12). On the other hand, it must be $L(\rho') \leq L(\rho)$ because of the assumption that the flow ρ' minimizes the maximum latency. Hence, it must be $L(\rho') = L(\rho)$ and $C(\rho') = C(\rho)$. In addition, since the function $C(\rho)$ is strictly convex and the set of feasible splittable flows forms a convex polytope, there is a unique flow which minimizes the total latency. Thus, ρ and ρ' must be identical. ■

Corollary 7.1 *A flow ρ minimizes the maximum latency if and only if for every $\pi_1, \pi_2 \in \mathcal{P}$ with $\rho(\pi_1) > 0$, $\theta_{\pi_1}(\rho) \leq \theta_{\pi_2}(\rho)$.*

Proof: By Lemma 7.1, the flow ρ minimizes the maximum latency if and only if it minimizes the total latency. Then, the corollary follows from the fact that ρ minimizes the total latency if and only if for every $\pi_1, \pi_2 \in \mathcal{P}$ with $\rho(\pi_1) > 0$, $\theta_{\pi_1}(\rho) \leq \theta_{\pi_2}(\rho)$ (eg, [30, Section 7.2], [32, Corollary 4.2]). ■

The following corollary states that the optimal splittable flow defines a mixed NE where all players adopt exactly the same strategy.

Corollary 7.2 *Let ρ_f^* be the optimal splittable flow and let \mathbf{p} be the mixed strategies profile where every player routes its traffic on each path π with probability $\rho_f^*(\pi)/W_{\text{tot}}$. Then, \mathbf{p} is a NE.*

Proof: By construction, the expected path loads corresponding to \mathbf{p} are equal to the values of ρ_f^* on these paths. Since all players follow exactly the same strategy and route their demand on each path π with probability ρ_f^*/W_{tot} , for each player i ,

$$\theta_{\pi}^{-i}(\mathbf{p}) = \theta_{\pi}(\mathbf{p}) - w_i \sum_{\pi' \in \mathcal{P}} |\pi \cap \pi'| \frac{\rho_f^*(\pi')}{W_{\text{tot}}} = (1 - \frac{w_i}{W_{\text{tot}}}) \theta_{\pi}(\mathbf{p})$$

Since the flow ρ_f^* also minimizes the total latency, for every $\pi_1, \pi_2 \in \mathcal{P}$ with $\rho_f^*(\pi_1) > 0$, $\theta_{\pi_1}(\mathbf{p}) \leq \theta_{\pi_2}(\mathbf{p})$ (eg, [6], [30, Section 7.2], [32, Corollary 4.2]), which also implies that $\theta_{\pi_1}^{-i}(\mathbf{p}) \leq \theta_{\pi_2}^{-i}(\mathbf{p})$. Therefore, for every player i and every $\pi_1, \pi_2 \in \mathcal{P}$ such that the player i routes its demand on π_1 with positive probability, $\lambda_{\pi_1}^i(\mathbf{p}) = \ell w_i + \theta_{\pi_1}^{-i}(\mathbf{p}) \leq \ell w_i + \theta_{\pi_2}^{-i}(\mathbf{p}) = \lambda_{\pi_2}^i(\mathbf{p})$. Consequently, \mathbf{p} is a NE. ■

7.4 An Upper Bound on the Social Cost.

Next we derive an upper bound on the social cost of every strategy profile whose maximum expected delay (ie, the maximum latency of its associated flow) is within a constant factor from the maximum latency of the optimal unsplittable flow.

Lemma 7.2 *Let ρ^* be the optimal unsplittable flow, and let \mathbf{p} be a mixed strategies profile and ρ its corresponding flow. If $L(\mathbf{p}) = L(\rho) \leq \alpha L(\rho^*)$, for some $\alpha \geq 1$, then*

$$\text{SC}(\mathbf{p}) \leq (\alpha + 1) \mathcal{O}(\frac{\log m}{\log \log m}) L(\rho^*),$$

where $m = |E|$ denotes the number of edges in the network.

Proof: For each edge $e \in E$ and each player i , let $X_{e,i}$ be the random variable describing the actual load routed through e by i . The random variable $X_{e,i}$ is equal to w_i if i routes its demand on a path π including e and 0 otherwise. Consequently, the expectation of $X_{e,i}$ is equal to $\mathbb{E}\{X_{e,i}\} = \sum_{\pi: e \in \pi} w_i p_i(\pi)$. Since each player selects its path independently, for every fixed edge e , the random variables $\{X_{e,i}, i \in [n]\}$, are independent from each other.

For each edge $e \in E$, let $X_e = \sum_{i=1}^n X_{e,i}$ be the random variable that describes the actual load routed through e , and thus, also the actual delay paid by any player traversing e . X_e is the sum of n independent random variables with values in $[0, w_{\max}]$. By linearity of expectation,

$$\mathbb{E}\{X_e\} = \sum_{i=1}^n \mathbb{E}\{X_{e,i}\} = \sum_{i=1}^n w_i \sum_{\pi \ni e} p_i(\pi) = \theta_e(\rho).$$

By applying the standard Hoeffding bound⁵ with $w = w_{\max}$ and $t = \exp \kappa \max\{\theta_e(\rho), w_{\max}\}$, we obtain that for every $\kappa \geq 1$,

$$\mathbb{P}\{X_e \geq \exp \kappa \max\{\theta_e(\rho), w_{\max}\}\} \leq \kappa^{-\exp \kappa}.$$

For $m \equiv |E|$, by applying the union bound we conclude that

$$\mathbb{P}\{\exists e \in E : X_e \geq \exp \kappa \max\{\theta_e(\rho), w_{\max}\}\} \leq m \kappa^{-\exp \kappa} \quad (13)$$

For each path $\pi \in \mathcal{P}$ with $\rho(\pi) > 0$, we define the random variable $X_\pi = \sum_{e \in \pi} X_e$ describing the actual delay along π . The social cost of \mathbf{p} , which is equal to the expected maximum delay experienced by some player, cannot exceed the expected maximum delay among paths π with $\rho(\pi) > 0$. Formally,

$$\text{SC}(\mathbf{p}) \leq \mathbb{E}\left\{\max_{\pi: \rho(\pi) > 0} \{X_\pi\}\right\}.$$

If for all $e \in E$, $X_e \leq \exp \kappa \max\{\theta_e(\rho), w_{\max}\}$, then for every path $\pi \in \mathcal{P}$ with $\rho(\pi) > 0$,

$$\begin{aligned} X_\pi = \sum_{e \in \pi} X_e &\leq \exp \kappa \sum_{e \in \pi} \max\{\theta_e(\rho), w_{\max}\} \\ &\leq \exp \kappa \sum_{e \in \pi} (\theta_e(\rho) + w_{\max}) \\ &= \exp \kappa (\theta_\pi(\rho) + \ell w_{\max}) \\ &\leq \exp \kappa (L(\rho) + \ell w_{\max}) \\ &\leq \exp (\alpha + 1) \kappa L(\rho^*) \end{aligned}$$

The third equality follows from $\theta_\pi(\rho) = \sum_{e \in \pi} \theta_e(\rho)$, the fourth inequality from $\theta_\pi(\rho) \leq L(\rho)$ since $\rho(\pi) > 0$, and the last inequality from the hypothesis that $L(\rho) \leq \alpha L(\rho^*)$ and the fact that $\ell w_{\max} \leq L(\rho^*)$ because ρ^* is an unsplittable flow. Therefore, using Ineq. (13), we conclude that

$$\mathbb{P}\left\{\max_{\pi: \rho(\pi) > 0} \{X_\pi\} \geq \exp (\alpha + 1) \kappa L(\rho^*)\right\} \leq m \kappa^{-\exp \kappa}.$$

⁵We use the standard version of Hoeffding bound ([19]): Let X_1, X_2, \dots, X_n be independent random variables with values in the interval $[0, w]$. Let $X = \sum_{i=1}^n X_i$ and let $\mathbb{E}\{X\}$ denote its expectation. Then, $\forall t > 0$, $\mathbb{P}\{X \geq t\} \leq \left(\frac{\exp \mathbb{E}\{X\}}{t}\right)^{t/w}$.

In other words, the probability that the actual maximum delay caused by \mathbf{p} exceeds the optimal maximum delay by a factor greater than $2 \exp(\alpha+1)\kappa$ is at most $m\kappa^{-\exp \kappa}$. Therefore, for every $\kappa_0 \geq 2$,

$$\begin{aligned} \text{SC}(\mathbf{p}) &\leq \mathbb{E} \left\{ \max_{\pi: \rho(\pi) > 0} \{X_\pi\} \right\} \leq \exp(\alpha+1)L(\rho^*) \left(\kappa_0 + \sum_{k=\kappa_0}^{\infty} kmk^{-\exp k} \right) \\ &\leq \exp(\alpha+1)L(\rho^*) (\kappa_0 + 2m\kappa_0^{-\exp \kappa_0+1}) \end{aligned}$$

If $\kappa_0 = \frac{2 \log m}{\log \log m}$, then $\kappa_0^{-\exp \kappa_0+1} \leq m^{-1}$, $\forall m \geq 4$. Thus, $\text{SC}(\mathbf{p}) \leq 2 \exp(\alpha+1) \left(\frac{\log m}{\log \log m} + 1 \right) L(\rho^*)$. \blacksquare

7.5 Bounding the Coordination Ratio.

Our final step is to show that the maximum expected delay of every NE is a good approximation to the optimal maximum latency. Then, we can apply Lemma 7.2 to bound the coordination ratio for our selfish routing game.

Lemma 7.3 *For every flow ρ corresponding to a mixed strategies profile \mathbf{p} at NE, $L(\rho) \leq 3L(\rho^*)$.*

Proof: The proof is based on Dorn's Theorem [10] which establishes strong duality in quadratic programming⁶. We use quadratic programming duality to prove that for any flow ρ at Nash equilibrium, the minimum expected delay $\theta^{\min}(\rho)$ cannot exceed $L(\rho_f^*) + \ell w_{\max}$. This implies the lemma because $L(\rho) \leq \theta^{\min}(\rho) + \ell w_{\max}$, since ρ is at Nash equilibrium, and $L(\rho^*) \geq \max\{L(\rho_f^*), \ell w_{\max}\}$, since ρ^* is an unsplittable flow.

Let Q be the square matrix describing the number of edges shared by each pair of paths. Formally, Q is a $|\mathcal{P}| \times |\mathcal{P}|$ matrix and for every $\pi, \pi' \in \mathcal{P}$, $Q[\pi, \pi'] = |\pi \cap \pi'|$. By definition, Q is symmetric. Next we prove that Q is positive semi-definite⁷.

$$\begin{aligned} x^T Q x &= \sum_{\pi \in \mathcal{P}} x(\pi) \sum_{\pi' \in \mathcal{P}} Q[\pi, \pi'] x(\pi') \\ &= \sum_{\pi \in \mathcal{P}} x(\pi) \sum_{\pi' \in \mathcal{P}} |\pi \cap \pi'| x(\pi') \\ &= \sum_{\pi \in \mathcal{P}} x(\pi) \sum_{e \in \pi} \sum_{\pi': e \in \pi'} x(\pi') \\ &= \sum_{\pi \in \mathcal{P}} x(\pi) \sum_{e \in \pi} \theta_e(x) \\ &= \sum_{e \in E} \theta_e(x) \sum_{\pi: e \in \pi} x(\pi) \\ &= \sum_{e \in E} \theta_e^2(x) \geq 0 \end{aligned}$$

⁶Let $\min\{x^T Q x + c^T x : Ax \geq b, x \geq \mathbf{0}\}$ be the primal quadratic program. The Dorn's dual of this program is $\max\{-y^T Q y + b^T y : A^T y - 2Qy \leq c, y \geq \mathbf{0}\}$. Dorn [10] proved strong duality when the matrix Q is symmetric and positive semi-definite. Thus, if Q is symmetric and positive semi-definite and both the primal and the dual programs are feasible, their optimal solutions have the same objective value.

⁷A $n \times n$ matrix Q is positive semi-definite if for every vector $x \in \mathbb{R}^n$, $x^T Q x \geq 0$.

First recall that for each edge e , $\theta_e(x) \equiv \sum_{\pi: e \in \pi} x(\pi)$. The third and the fifth equalities follow by reversing the order of summation. In particular, in the third equality, instead of considering the edges shared by π and π' , for all $\pi' \in \mathcal{P}$, we consider all the paths π' using each edge $e \in \pi$. On both sides of the fifth inequality, for every edge $e \in E$, $\theta_e(x)$ is multiplied by the sum of $x(\pi)$ over all the paths π using e .

Let ρ also denote the $|\mathcal{P}|$ -dimensional vector corresponding to the flow ρ . Then, the π -th coordinate of $Q\rho$ is equal to the expected delay $\theta_\pi(\rho)$ on the path π , and the total latency of ρ is $C(\rho) = \rho^T Q\rho$.

Therefore, the problem of computing a feasible splittable flow of minimum total latency is equivalent to computing the optimal solution to the following quadratic program: $\min\{\rho^T Q\rho : \mathbf{1}^T \rho \geq W_{\text{tot}}, \rho \geq \mathbf{0}\}$, where $\mathbf{1}/\mathbf{0}$ denotes the $|\mathcal{P}|$ -dimensional vector having 1/0 in each coordinate. Also notice that no flow of value strictly greater than W_{tot} can be optimal for this program. This quadratic program is clearly feasible and its optimal solution is ρ_f^* (Lemma 7.1).

The Dorn's dual of this quadratic program is: $\max\{zW_{\text{tot}} - \rho^T Q\rho : 2Q\rho \geq \mathbf{1}z, z \geq 0\}$ (e.g., [10], [5, Chapter 6]). We observe that any flow ρ can be regarded as a feasible solution to the dual program by setting $z = 2\theta^{\min}(\rho)$. Hence, both the primal and the dual programs are feasible. By Dorn's Theorem [10], the objective value of the optimal dual solution is exactly $C(\rho_f^*)^8$.

Let ρ be any feasible flow at Nash equilibrium. Setting $z = 2\theta^{\min}(\rho)$, we obtain a dual feasible solution. By the discussion above, the objective value of the feasible dual solution $(\rho, 2\theta^{\min}(\rho))$ cannot exceed $C(\rho_f^*)$. In other words,

$$2\theta^{\min}(\rho)W_{\text{tot}} - C(\rho) \leq C(\rho_f^*) \quad (14)$$

Since ρ is at Nash equilibrium, $L(\rho) \leq \theta^{\min}(\rho) + \ell w_{\max}$. In addition, by Ineq. (11), the average latency of ρ cannot exceed its maximum latency. Thus,

$$C(\rho) \leq L(\rho)W_{\text{tot}} \leq \theta^{\min}(\rho)W_{\text{tot}} + \ell w_{\max}W_{\text{tot}}$$

Combining the inequality above with Ineq. (14), we obtain that $\theta^{\min}(\rho)W_{\text{tot}} \leq C(\rho_f^*) + \ell w_{\max}W_{\text{tot}}$. Using $C(\rho_f^*) = L(\rho_f^*)W_{\text{tot}}$, we conclude that $\theta^{\min}(\rho) \leq L(\rho_f^*) + \ell w_{\max}$. ■

The following theorem is an immediate consequence of Lemma 7.3 and Lemma 7.2.

Theorem 8 ([15]) *The price of anarchy of any ℓ -layered network congestion game with resource delays equal to their loads, is at most $8 \exp(\frac{\log m}{\log \log m} + 1)$.*

A recent development which is complementary to the last theorem is the following which we state without a proof:

Theorem 9 ([16]) *The price of anarchy of any unweighted, single commodity network congestion game with resource delays $(d_e(x) = a_e \cdot x, a_e \geq 0)_{e \in E}$, is at most $24 \exp(\frac{\log m}{\log \log m} + 1)$.*

⁸More specifically, the optimal dual solution is obtained from ρ_f^* by setting $z = 2\theta^{\min}(\rho_f^*)$. Since $L(\rho_f^*) = \theta^{\min}(\rho_f^*)$ and $C(\rho_f^*) = L(\rho_f^*)W_{\text{tot}}$, the objective value of this solution is $2\theta^{\min}(\rho_f^*)W_{\text{tot}} - C(\rho_f^*) = C(\rho_f^*)$.

8 The Pure Price of Anarchy in Congestion Games

In this last section we overview some recent advances in the *Pure Price of Anarchy* (PPoA) of congestion games, that is, the worst-case ratio of the social cost of a PNE over the social optimum of the game.

The case of linear resource delays has been extensively studied in the literature. The PPoA wrt the total latency objective has been proved that it is $\frac{3+\sqrt{5}}{2}$, even for weighted multicommodity network congestion games [4, 7]. This result is also extended to the case of mixed equilibria. For the special case of identical players it has been proved (independently by the papers [4, 7]) that the PPoA drops down to $5/2$. When considering identical users and single-commodity network congestion games, the PPoA is again $5/2$ wrt the maximum latency objective, but explodes to $\Theta(\sqrt{n})$ for multicommodity network congestion games ([7]). Earlier it was implicitly proved by [15] that the PPoA of any weighted congestion game on a layered network with resource delays identical to the congestion, is at most 3.

9 Conclusions

In this survey we have presented some of the most significant advances concerning the atomic (mainly network) congestion games literature. We have focused on issues dealing with existence of PNE, construction of an arbitrary PNE when such equilibria exist, as well as the price of anarchy for many broad subclasses of network congestion games.

We highlighted the significance of allowing distinguishable players (ie, players with different action sets, or with different traffic demands, or both) and established some kind of “equivalence” between games with unit-demand players on arbitrary networks with delays equal to the loads and games with players of varying demands on layered networks.

Still, there remain many unresolved questions. The most important question is a complete characterization of the subclasses of games having PNE and admitting polynomial time algorithms for constructing them, in the case of general networks.

Additionally, a rather recent trend deals with the network-design perspective of congestion games. For these games [3, 2, 20] the measure of performance is the **price of stability**, ie, the ratio of the *best* NE over the social optimum of the game, trying to capture the notion of the gap between solutions proposed by the network designer to the players and the social optimum of the game (which may be an unstable state for the players). This seems to be a rather intriguing and very interesting (complementary to the one presented here) approach of congestion games, in which there are still many open questions.

References

- [1] Ahuja R., Magnanti T., Orlin J. *Network Flows: Theory, Algorithms and Applications*. Prentice Hall, 1993.
- [2] Anshelevich E., Dasgupta A., Kleinberg J., Tardos E., Wexler T., Roughgarden T. 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)*, page to appear, 295–304.

- [3] Anshelevich E., Dasgupta A., Tardos E., Wexler T. Near-optimal network design with selfish agents. In *Proc. of the 35th ACM Symp. on Theory of Computing (STOC '03)*, pages 511–520, 2004.
- [4] Awerbuch B., Azar Y., Epstein A. The price of routing unsplittable flow. In *Proc. of the 37th ACM Symp. on Theory of Computing (STOC '05)*, page to appear, 2005.
- [5] Bazaraa M.S., Sherali H.D., Shetty C.M. *Nonlinear Programming: Theory and Algorithms (2nd edition)*. John Wiley and Sons, Inc., 1993.
- [6] Beckmann M., McGuire C.B., Winsten C.B. *Studies in the Economics of Transportation*. Yale University Press, 1956.
- [7] G. Christodoulou and E. Koutsoupias. The Price of Anarchy of Finite Congestion Games. In *stoc05*, pages to appear, available from <http://cgi.di.uoa.gr/~elias/publications/paper--ck05.ps>, 2005.
- [8] Czumaj A., Krysta P., Voecking B. Selfish traffic allocation for server farms. In *Proc. of the 34th ACM Symp. on Theory of Computing (STOC '02)*, pages 287–296, 2002.
- [9] Czumaj A., Voecking B. Tight bounds for worst-case equilibria. In *Proc. of the 13th ACM-SIAM Symposium on Discrete Algorithms (SODA '02)*, pages 413–420, 2002.
- [10] Dorn W.S. Duality in quadratic programming. *Quarterly of Applied Mathematics*, 18(2):155–162, 1960.
- [11] Even-Dar E., Kesselman A., Mansour Y. Convergence time to nash equilibria. In *Proc. of the 30th International Colloquium on Automata, Languages and Programming (ICALP '03)*, pages 502–513. Springer-Verlag, 2003.
- [12] Fabrikant A., Papadimitriou C., Talwar K. The complexity of pure nash equilibria. In *Proc. of the 36th ACM Symp. on Theory of Computing (STOC '04)*, 2004.
- [13] Feldmann R., Gairing M., Luecking T., Monien B., Rode M. Nashification and the coordination ratio for a selfish routing game. In *Proc. of the 30th International Colloquium on Automata, Languages and Programming (ICALP '03)*, pages 514–526. Springer-Verlag, 2003.
- [14] Fotakis D., Kontogiannis S., Koutsoupias E., Mavronicolas M., Spirakis P. The structure and complexity of nash equilibria for a selfish routing game. In *Proc. of the 29th International Colloquium on Automata, Languages and Programming (ICALP '02)*, pages 123–134. Springer-Verlag, 2002.
- [15] Fotakis D., Kontogiannis S., Spirakis P. Selfish unsplittable flows. In *Proc. of the 31st International Colloquium on Automata, Languages and Programming (ICALP '04)*, pages 593–605. Springer-Verlag, 2004.
- [16] Fotakis D., Kontogiannis S., Spirakis P. Symmetry in network congestion games: Pure equilibria and anarchy cost. In *in progress (submitted for publication)*, 2005.
- [17] Gairing M., Luecking T., Mavronicolas M., Monien B., Spirakis P. Extreme nash equilibria. In *8th Italian Conference on Theoretical Computer Science (ICTCS'03)*. Springer-Verlag, 2003.

- [18] Hochbaum D., Shmoys D. A polynomial approximation scheme for scheduling on uniform processors: Using the dual approximation approach. *SIAM J. Comput.*, 17(3):539–551, 1988.
- [19] Hoeffding W. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- [20] Johari R., Tsitsiklis J. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, page to appear, 2004.
- [21] Johnson D., Papadimitriou C., Yannakakis M. How easy is local search? *J. of Computer and System Sciences*, 37:79–100, 1988.
- [22] Koutsoupias E., Mavronicolas M., Spirakis P. Approximate equilibria and ball fusion. *ACM Transactions on Computer Systems*, 36:683–693, 2003.
- [23] Koutsoupias E., Papadimitriou C. Worst-case equilibria. In *Proc. of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS '99)*, pages 404–413. Springer-Verlag, 1999.
- [24] Libman L., Orda A. Atomic resource sharing in noncooperative networks. *Telecommunication Systems*, 17(4):385–409, 2001.
- [25] Lücking T., Mavronicolas M., Monien B., Rode M. A new model for selfish routing. In *Proc. of the 21st Annual Symposium on Theoretical Aspects of Computer Science (STACS '04)*, page to appear. Springer-Verlag, 2004.
- [26] Lücking T., Mavronicolas M., Monien B., Rode M., Spirakis P., Vrto I. Which is the worst-case nash equilibrium? In *26th International Symposium on Mathematical Foundations of Computer Science (MFCS'03)*, pages 551–561. Springer-Verlag, 2003.
- [27] Mavronicolas M., Spirakis P. The price of selfish routing. In *Proc. of the 33rd ACM Symp. on Theory of Computing (STOC '01)*, pages 510–519, 2001.
- [28] Milchtaich I. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13:111–124, 1996.
- [29] Monderer D., Shapley L. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
- [30] Papadimitriou C., Steiglitz K. *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall, Inc., 1982.
- [31] Rosenthal R.W. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
- [32] Roughgarden T., Tardos E. How bad is selfish routing? *J. Assoc. Comput. Mach.*, 49(2):236–259, 2002.
- [33] Roughgarden T. The price of anarchy is independent of the network topology. In *Proc. of the 34th ACM Symp. on Theory of Computing (STOC '02)*, pages 428–437, 2002.
- [34] Schäffer A., Yannakakis M. Simple local search problems that are hard to solve. *SIAM J. on Computing*, 20(1):36–87, 1991.