Atomic Congestion Games among Coalitions *

Dimitris Fotakis[1], Spyros Kontogiannis[2,3], and Paul Spirakis[3]

[1] Dept. of Information and Communication Systems Engineering,
University of the Aegean, 83200 Samos, Greece. fotakis@aegean.gr

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

[3] Research Academic Computer Technology Institute,
P.O. Box 1382, N. Kazantzaki Str., 26500 Rio-Patra, Greece. {kontog, spirakis}@cti.gr

June 21, 2006

Abstract. We consider algorithmic questions concerning the existence, tractability and quality of Nash equilibria in atomic congestion games among users participating in (static) selfish coalitions.
We introduce a coalitional congestion model among atomic players and demonstrate many interesting similarities with the non-cooperative case. For example, there exists a potential function proving the existence of Pure Nash Equilibria (PNE) in the unrelated parallel links setting; in the network setting, the Finite Improvement Property collapses as soon as we depart from linear delays, but there is an exact potential (and thus PNE) for linear delays; the Price of Anarchy on identical parallel links demonstrates a quite surprising threshold behavior: it persists on being asymptotically equal to that in the case of the non-cooperative KP-model, unless the number of coalition is sublogarithmic.
We also show crucial differences, mainly concerning the hardness of algorithmic problems that are solved efficiently in the non-cooperative case. Although we demonstrate convergence to robust PNE, we also prove the hardness of computing them. On the other hand, we can easily construct a generalized fully mixed Nash Equilibrium. Finally, we propose two natural improvement policies and show that both converging in pseudo-polynomial time to PNE that are robust against (even dynamically forming) coalitions of small size.

Keywords. Game Theory, Atomic Congestion Games, Coalitions, Convergence to Equilibria, Price of Anarchy.

1 Introduction

The new research field of algorithmic game theory, till now, had its focus mostly on *non–cooperative* strategic games. Indeed, the cases where dynamic game elements are studied are rather rare. In addition, the main concept of algorithmic game theory, namely, the *price of anarchy* (PoA), has been applied mostly to non–cooperative selfish players with a variety of pure strategy sets and payoff functions. Most of the examples are motivated from network traffic and congestion problems. But real life examples justify the necessity for the consideration of selfish coalitions, since in most cases there is some sort of hierarchy that needs to be taken into account. For example, in a communication network, the service providers are interested in minimizing their own cost for assuring a promised quality of service (eg, bandwidth) to their users, but they are not actually interested in utilizing their users' individual delays, so long as the maximum delay is small. In fact, in some cases they have to sacrifice the utilization of some of their users dictatorially, for the sake of their own (still private and selfish wrt

* A preliminary version of this work appeared in the Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP 2006), Lecture Notes in Computer Science 4051. This work was partially supported by the EU within 6th Framework Programme under contract 001907 (DELIS).

other coalitions) objective. In this scenario, each service provider can be seen as a static coalition of users that tries to minimize the maximum cost that any of its users has to pay. Alternatively, we may actually have altruistic (wrt their participating users) coalitions that try to minimize the cumulative cost that all their users have to pay, by performing joint decisions for all their members.

All these scenarios are captured by (static) coalitional congestion games, in which the selfish players are actually the coalitions that may handle more than one users at will. In this work we are not interested in how this coalitional cost is shared among the members of the coalition, neither for the dynamics that lead to coalition formation. Our main goal is to study the existence, convergence, and quality of produced Nash Equilibria among given (static) coalitions, that are robust not only against unilateral moves of the users, but also against joint moves of subsets of any coalition.

Initially we propose a framework for studying these games and consequently we provide some evidence on the difficulty of classical algorithmic questions that mainly concern strategic games, such as the existence and tractability of Nash Equilibria (for the coalitions), or estimating the quality of these games.

1.1 Related Work

Existence of PNE and convergence issues have been extensively studied for the non–cooperative KP Model (eg, [8, 9, 11, 12, 21, 26] and references therein) and networks (eg, [13, 14, 10, 21, 26]). It is well known that there is always a PNE in the non–cooperative KP Model, even for the case of unrelated parallel machines. It is quite impressive that convergence to PNE in the non–cooperative KP Model may vary dramatically, depending on the families of improvement paths we consider. For example, in the case of identical links it may take an exponential number of steps for some family (SHORTEST JOB FIRST) and linear number of steps for another (LONGEST JOB FIRST) [8] to converge to some PNE. A quite interesting perspective is that of selfish rerouting in parallel, proposed by [9]. This is an issue that is not dealt with in the present paper, and seems to be very promising direction for future research in coalitional congestion games.

The quality of atomic non–cooperative congestion games has also been a hot topic in the recent literature, both for arbitrary Nash Equilibria (eg, [6, 12–14, 20, 22] and references therein) and only among pure Nash Equilibria (eg, [1, 2]). For an overview of recent developments on algorithmic questions on non–cooperative atomic congestion games, the interested reader is referred to [18].

In this work we focus on coalitions of players in atomic congestion games. This implies that each coalition selfishly governs a unique subset of atomic players, whose traffic demands have to be routed via single paths (ie, unsplittably). Nevertheless, a coalition is allowed to choose different routes for different players that it handles. Our concern is the existence and construction of Pure Nash Equilibria in these games, as well as the effect of coalition formation to the quality of the game.

A similar concept was discussed by Orda et al [26], Roughgarden [27], and Correa et al. [5] in a different framework (called *atomic splittable* setting), where each coalition has its own atomic weight which it routes selfishly from a source to a destination node in the network, but is allowed to split this weight into infinitesimally small pieces. Orda et al. [26] show that in a single-commodity network congestion game with coalitions and the selfish cost being the total delay of the coalition, NE may not be unique even for two coalitions of different cardinalities. On the other hand, the uniqueness of NE is established for the special case of parallel links. Recently, Cominetti et al. [3] proved that the proofs in [5, 27] have a basic flaw and many of the provided bounds on the price of anarchy are actually incorrect. They also provided the correct (albeit slightly weaker) bounds in some cases.

Very recently (and independently), Hayrapetyan et al. [15] have dealt with the effect of coalition formation in congestion games. They consider static coalitions as well, but they focus on the objective of *total delay* (each coalition has to pay for the sum of delays that its users suffer). They consider coalitions among splittable flows (which is closely related to the atomic splittable congestion games studied in [26, 5, 27]), as well as coalitions among atomic players (ie, whose traffic demands cannot be split at all). But they only consider the case where each atomic player has unit traffic demand. They prove that coalition formation is not necessarily beneficial to the quality of the game (compared to the coalition-free game), unless we have an atomic splittable congestion game in a network of parallel links with convex link delay functions.

In our work, we mainly focus on coalitions among atomic users with arbitrary (integer) traffic demands, that have to assure a certain quality of service for the users, but are otherwise completely selfish. Our main concern is the coalitional KP Model. Therefore, we consider the objective of *maximum delay*, where each coalition has to pay for the maximum delay among its users. In our case, it is easy to show that coalitions are always beneficial (or at least, not worse) for the quality of the game. We also provide some preliminary results for coalitions in weighted congestion games.

1.2 Our Contribution and Roadmap

In section 2 we define the coalitional congestion models that we study and provide the necessary definitions and notation. The emphasis is on parallel links with a coalition's selfish cost being the maximum delay that any of its users suffers (aka the KP Model [20]).

In section 3, we prove the existence of PNE in the coalitional KP Model, even for dynamically forming coalitions and unrelated parallel links. In section 4, we focus on identical parallel links and establish pseudo-polynomial convergence to 2-robust PNE for two natural families of improvement paths. The first family is called SMALLER COALITIONS FIRST and gives priority to selfish moves of smaller coalitions that involve the smallest number of links possible. The analysis highlights the importance of a special class of selfish 2-moves called $2-$flips. In the second family called BEST OF THE BEST RESPONSE, every coalition switches to its best response that also minimizes the weighted total delay of its users. We conjecture that these results can be generalized to k-robust PNE, for any constant $k \geqslant 2$.

For efficient constructions of Nash Equilibria we have both bad and good news. We prove that computing a PNE for coalitional congestion games on identical parallel link remains $\mathbf{NP}$-complete even if there are almost as many coalitions as users. On the positive side, we show pseudo-polynomial convergence time for the SMALLER COALITIONS FIRST and BEST OF THE BEST RESPONSE dynamics. In addition, we show that any coalitional congestion game on identical parallel links admits a generalized fully mixed NE.

In section 5 we prove asymptotically tight bounds for the Price of Anarchy of the Coalitional KP Model (with identical links). Indeed, we demonstrate a *threshold behavior* of the Price of Anarchy wrt the number of coalitions: It is as bad as in the case of the traditional (coalition-free) KP Model, so long as there are $\Omega(\frac{\log m}{\log\log m})$ coalitions. From that point on, the Price of Anarchy drops linearly with the number of coalitions.

Finally, in section 6 we deal with weighted network congestion games that allow static coalitions of players. A coalition's selfish cost is the weighted total delay of its users. We prove that even very simple unweighted single-commodity network congestion games with linear or 2-wise linear delay functions may not possess the Finite Improvement Property, although a PNE may exist. Therefore, we cannot hope for a potential function of any kind for these games. On the other hand, if we restrict our attention to linear delays, every weighted congestion game admits an exact potential identical to that of the non-cooperative case [13].

2 The Model of Static Coalitional Congestion Games

In this section we depart from the classical non-cooperative approach of congestion games and consider the case where the users are allowed to form coalitions that behave as single players of the game. Each coalitional player attempts to maximize its own private objective, given the strategies of the other coalitional players; it does not actually care about the private costs of the individual users whose coalition it represents. We initially consider the case of parallel links (the classical KP model) and consequently we deal with networks.

2.1 Coalitional KP Model

We consider a collection $[m]$ of identical parallel machines (links) and a collection $[n]$ of jobs[1]. Each of these jobs may be allocated to any of the m available links. Associated with each job $i \in [n]$ is an integer *service demand* $w_i \in \mathbb{N}_+$ (eg, the number of elementary operations for the execution of job i). Denote by $\widetilde{W} = \{w_i\}_{i\in[n]}$ the multiset of the jobs' demands.

Definition 1 (Coalitions). *A set of $k \geqslant 1$ (static) coalitions $C_1, \ldots, C_k$ is a fixed partition of $\widetilde{W}$ into k nonempty multisets. Hence: (i) the union (as multisets) of these coalitions is exactly $\widetilde{W}$, (ii) $C_j \neq \emptyset$, $\forall j \in [k]$, and (iii) $C_i \cap C_j = \emptyset$, $\forall i, j \in [k] : i \neq j$.*

For $j \in [k]$ let $C_j = \{w_j^1, \ldots, w_j^{n_j}\}$, so that $\sum_{j=1}^{k} n_j = n$. Let $W_j = \sum_{i=1}^{n_j} w_j^i$ be the cumulative service demand required by coalition C_j, let $w_j^{\max} = \max_{i\in[n_j]}\{w_j^i\}$ be the largest task of coalition C_j, and let $W_{\text{tot}} = \sum_{j\in[k]} W_j$ be the overall service demand required by the system. Wlog assume that $w_j^1 \geqslant \cdots \geqslant w_j^{n_j}$ for all $j \in [k]$.

Strategies and Profiles. A *pure strategy* $\sigma_j = (\sigma_j^i)_{i\in[n_j]}$ for coalition C_j defines the deterministic selection of a link $\sigma_j^i \in M$ for each $w_j^i \in C_j$. Denote by Σ_j the set of all pure strategies available to coalition C_j. Clearly, $\Sigma_j = M^{n_j}$. We let $S_j(\sigma_j) = \{\ell \in [m] : \exists i \in [n_j] \text{ with } \sigma_j^i = \ell\}$ denote the set of links getting some task of coalition j in pure strategy σ_j.

A *mixed strategy* for coalition C_j is a probability distribution $\mathbf{p_j}$ on the set Σ_j of its pure strategies (ie, a point of the simplex $\Delta(\Sigma_j) \equiv \{\mathbf{q} \in \mathbb{R}^{n_j} : \mathbf{q} \geqslant \mathbf{0};\ \mathbf{1}^T\mathbf{q} = 1\}$). In order to indicate the probability of pure strategy σ_j being chosen by C_j when the mixed strategy $\mathbf{p_j}$ has been adopted, we use (for sake of simplicity) the functional notation $p_j(\sigma_j)$, rather than the coordinate of the vector $\mathbf{p_j}$ corresponding to σ_j.

A *pure strategies profile* or *configuration* for the coalitions is a collection $\sigma = (\sigma_j)_{j\in[k]}$ of pure strategies, one per coalition. $\Sigma \equiv \times_{j\in[k]}\Sigma_j$ is the set of all the possible configurations of the game (called the *configuration space*). Let (σ_{-j}, α_j) denote the configuration resulting from a configuration σ when coalition C_j unilaterally changes its pure strategy from σ_j to α_j. Similarly, we call the simplotope $\Delta(\Sigma) \equiv \times_{j\in[k]}\Delta(\Sigma_j)$ the *mixed strategies space* of the coalitional game. A *mixed strategies profile* $\mathbf{p} = (\mathbf{p_j})_{j\in[k]} \in \Delta(\Sigma)$ is a collection of mixed strategies, one per coalition. If for every coalition $j \in [k]$ and every configuration $\sigma_j \in \Sigma_j, p_j(\sigma_j) > 0$, $\mathbf{p}$ is a *fully mixed* profile.

A special case of particular interest in scheduling literature, is when the coalitions are enforced to eventually choose consecutive links for their own tasks (or similarly, in a single subinterval in case of interval scheduling). When the coalitions are forced to choose only this kind of pure strategies for their own tasks, then we shall refer to the *Coalitional Chains* model.

[1] For any integer $k \geqslant 1$, $[k] \equiv \{1, \ldots, k\}$.

Selfish Costs. Fix a configuration $\sigma = (\sigma_j)_{j \in [k]}$. Define the load on link $\ell \in M$ due to coalition C_j as $\theta_\ell(\sigma_j) \equiv \sum_{i \in [n_j]:\sigma_j^i = \ell} w_j^i$. Then the total load on link $\ell \in M$ is the total demand on link ℓ with respect to σ, ie, $\theta_\ell(\sigma) = \sum_{j=1}^{k} \theta_\ell(\sigma_j)$. Similarly, the load induced on link $\ell \in M$ by all the coalitions, except for coalition C_j is $\theta_\ell(\sigma_{-j}) = \sum_{j' \in [k] \setminus \{j\}} \theta_\ell(\sigma_{j'})$ The *selfish cost* $\lambda_j(\sigma)$ of coalition C_j is the maximum load over the set of machines it uses, formally $\lambda_j(\sigma) = \max_{\ell \in S_j(\sigma_j)} \{\theta_\ell(\sigma)\}$.

For a mixed profile $\mathbf{p}$, the load on each link $\ell \in M$ becomes a random variable induced by the probability distributions $\mathbf{p_j}$ for all $j \in [k]$. More precisely, let $\theta_\ell(\mathbf{p_j}) = \sum_{\sigma_j \in \Sigma_j} p_j(\sigma_j)\theta_\ell(\sigma_j)$ be the expected load induced on link ℓ by coalition j according to mixed strategy $\mathbf{p_j}$. The *expected load* on link $\ell \in M$, denoted $\theta_\ell(\mathbf{p})$, is the expectation of the load on ℓ according to $\mathbf{p}$. Formally,

$$\theta_\ell(\mathbf{p}) = \sum_{\sigma \in \Sigma} \left[\left(\prod_{j \in [k]} p_j(\sigma_j) \right) \cdot \theta_\ell(\sigma) \right] = \sum_{j \in [k]} \theta_\ell(\mathbf{p_j})$$

We let $\theta_\ell(\mathbf{p_{-j}})$ denote the expected load that all the coalitions except for coalition C_j induce on link ℓ. Formally,

$$\theta_\ell(\mathbf{p_{-j}}) = \sum_{\sigma_{-j} \in \Sigma_{-j}} \left[\left(\prod_{j' \in [k] \setminus \{j\}} p_{j'}(\sigma_{j'}) \right) \cdot \theta_\ell(\sigma_{-j}) \right] = \sum_{j' \in [k] \setminus \{j\}} \theta_\ell(\mathbf{p_{j'}})$$

The *conditional expected selfish cost* of coalition j adopting the pure strategy $\sigma_j \in \Sigma_j$, given that the other players follow the strategies indicated by $\mathbf{p}$, is

$$\lambda_j(\mathbf{p}, \sigma_j) = \max_{\ell \in S_j(\sigma_j)} \{\theta_\ell(\sigma_j) + \theta_\ell(\mathbf{p_{-j}})\}$$

In words, coalition C_j pays for the conditional expectation of its selfish cost, had it adopted the pure strategy $\sigma_j \in \Sigma_j$. This is because coalition C_j has to encounter all the possible alternatives for its own jobs prior to the other coalitions' determination of their actual action, knowing only their probability distributions. The *expected selfish cost* of coalition C_j is defined as the expectation of coalition C_j's conditional expected cost, over all possible actions that can be taken: $\lambda_j(\mathbf{p}) = \sum_{\sigma_j \in \Sigma_j} [p_j(\sigma_j) \cdot \lambda_j(\mathbf{p}, \sigma_j)]$.

Remark. Observe that each coalition pays for the *expected maximum load* that it would cause as if it was on its own, plus the *expected loads* caused by the other coalitions to each of the links.

Nash Equilibria. The definition of expected selfish costs completes the definition of the finite non-cooperative game involving the k coalitions of jobs that are to be assigned to the m links: $\Gamma = \langle [k], (\Sigma_j)_{j \in [k]}, (\lambda_j)_{j \in [k]} \rangle$. We are interested in the induced Nash Equilibria [25] of Γ. Informally, a Nash Equilibrium is a (pure or mixed) profile such that no coalition can reduce its selfish cost by unilaterally changing its strategy. Formally:

Definition 2 (Nash Equilibrium). *A pure strategies profile* $\sigma = (\sigma_j)_{j \in [k]}$ *is a* Pure Nash Equilibrium (PNE) *for* Γ *if,* $\forall j \in [k], \quad \forall \alpha_j \in \Sigma_j, \quad \lambda_j(\sigma) \leqslant \lambda_j(\sigma_{-j}, \alpha_j)$. *A mixed strategies profile* $\mathbf{p}$ *is a* Nash Equilibrium (NE) *if,* $\forall j \in [k]$, $\forall \sigma_j \in \Sigma_j$, *it holds that* $p_j(\sigma_j) > 0 \quad \Rightarrow \quad \sigma_j \in \arg\min_{\alpha_j \in \Sigma_j} \{\lambda_j(\mathbf{p}, \alpha_j)\}$.

Social Cost, Social Optimum and Price of Anarchy. For any configuration $\sigma = (\sigma_j)_{j\in[k]}$ we define the *social cost*, denoted $\mathsf{SC}(\sigma)$, to be the maximum load over the set of machines M with respect to σ, ie, $\mathsf{SC}(\sigma) = \max_{\ell\in M}\{\theta_\ell(\sigma)\} = \max_{j\in[k]}\{\lambda_j(\sigma)\}$. For any mixed profile $\mathbf{p}$ the social cost is defined as the expectation, over all random choices of the coalitions, of the maximum load over the set of machines: $\mathsf{SC}(\mathbf{p}) = \sum_{\sigma\in\Sigma}\left(\prod_{j=1}^{k} p_j(\sigma_j)\right)\cdot \max_{\ell\in M}\{\theta_\ell(\sigma)\}$.

Now let σ^* be a configuration that minimizes the social cost function, ie, $\sigma^* \in \arg\min_\sigma\{\mathsf{SC}(\sigma)\}$. Thus σ^* is an optimal configuration of the set of loads $\widetilde{W}$ to the set of machines M. We denote this value by $\mathsf{OPT} = \mathsf{SC}(\sigma^*)$. The *Price of Anarchy* (also referred to as *Coordination Ratio*) [20], is the worst–case ratio of the social cost paid at any NE over the value of the social optimum of the game: $\mathsf{R} = \max_{\mathbf{p}\text{ is NE}}\left\{\frac{\mathsf{SC}(\mathbf{p})}{\mathsf{OPT}}\right\}$.

Improvement Paths. When we discuss convergence issues, we shall frequently refer to the notion of improvement paths: These are sequences of configurations for the coalitions (ie, points in Σ), such that any two consecutive configurations differ only in the pure strategy of *exactly one* coalition, and additionally the cost of this unique coalition is *strictly less* in the latter configuration than in the former one.

2.2 Coalitional Players in Networks

Let $G = (V, E)$ be a directed network with a non-decreasing delay function $d_e(x)$ associated with each $e \in E$. We are also given a multiset of users $\widetilde{W} = \{w_i\}_{i\in[n]}$ each willing to route his traffic from a source to a destination. The congestion game is *unweighted* if all users have identical (unit) weights, and *weighted* otherwise. In a network congestion game with coalitions, the set of players is the set of coalitions $\{C_1, \ldots, C_k\}$. A coalition $C_j = \{w_j^1, \ldots, w_j^{n_j}\}$ consists of n_j users willing routing their traffic from s_j to t_j.

Let $\mathcal{P}_j$ be the set of $s_j - t_j$ paths in G. The set of pure strategies of coalition C_j is $\mathcal{P}_j^{n_j}$. Thus a pure strategy σ_j determines an $s_j - t_j$ path σ_j^i for every user $i \in [n_j]$. For every edge $e \in E$, the load of e due to the users of C_j in σ_j is $\theta_e(\sigma_j) = \sum_{i\in[n_j]:e\in\sigma_j^i} w_j^i$.

A pure strategies profile or configuration for the coalitions is a collection $\sigma = (\sigma_j)_{j\in[k]}$ of pure strategies, one per coalition. For every edge $e \in E$, the load of e in σ is $\theta_e(\sigma) = \sum_{j=1}^{k}\theta_e(\sigma_j)$. For every path $\pi \in \bigcup_{j=1}^{k}\mathcal{P}_j$, the delay along π in σ is $d_\pi(\sigma) = \sum_{e\in\pi} d_e(\theta_e(\sigma))$.

There are (at least) two natural notions of selfish cost of coalition C_j in σ. The first is the *maximum delay* over all paths used by C_j, denoted $\lambda_j(\sigma)$. Formally, $\lambda_j(\sigma) = \max_{i\in C_j}\{d_{\sigma_j^i}(\sigma)\}$. The second is the *weighted total delay* of coalition C_j, denoted $\tau_j(\sigma)$. Formally,

$$\tau_j(\sigma) = \sum_{i\in C_j} d_{\sigma_j^i}(\sigma) = \sum_{e\in E}\theta_e(\sigma_j)d_e(\theta_e(\sigma))$$

For networks other than parallel links, we use the weighted total delay as the notion of selfish cost.

A configuration σ is a Pure Nash Equilibrium (PNE) if no coalition can reduce its selfish cost by unilaterally rerouting its users. Formally, σ is a PNE if for every coalition C_j, and every pure strategy $\alpha_j \in \mathcal{P}_j^{n_j}$, $\tau_j(\sigma) \leqslant \tau_j(\sigma_{-j}, \alpha_j)$.

Remark. Both maximum delay and weighted total delay generalize naturally the notion of selfish cost used in weighted congestion games. Also note that computing a coalition's best response for total delay in unweighted single-commodity network congestion games can be performed by first applying the transformation of [10, Theorem 2] and then computing a min-cost flow. On the other hand, computing a coalition's best response for maximum delay is $\mathbf{NP}$-hard even for unweighted single-commodity network congestion games with linear delays and coalitions of size 2 (see e.g. [4, Theorem 3]).

3 Tractability of NE in the Coalitional KP Model

It is not hard to prove the existence of PNEs in coalitional congestion games, even if we assume unrelated parallel links and allow the players to form *dynamic coalitions* (ie, consider the case of different combinations of players that attempt a joint selfish move). This is done by showing the existence of a generalized ordinal potential function which assures the convergence to a PNE in a finite number of steps.

For the case of unrelated parallel links, we prove that any improvement path that combines selfish movements of (even dynamically forming) coalitions of size at most k, has length at most $(k+1)^{W_{\rm tot}}/k$, where for the general case of unrelated parallel links, $w_j(\ell) > 0$ denotes the weight of job j on link ℓ and $W_{\rm tot} = \sum_{j\in[n]} \max_{\ell\in[m]}\{w_j(\ell)\}$.

Theorem 1. *For the case of unrelated parallel links and selfish tasks with integer weights, any improvement path that combines arbitrary selfish movements of coalitions of players of size at most k (even if these are formed dynamically) has length at most $(k+1)^{W_{\rm tot}}/k$.*

Proof. Assume an arbitrary improvement path $\langle \varpi(0), \ldots, \varpi(T)\rangle$, where $\varpi(t) \in [m]^n$, $\forall 0 \leqslant t \leqslant T$. Assume any $1 \leqslant t \leqslant T$ and suppose that at time step t a coalition C of $k \geqslant 1$ players choose to cooperate and perform a selfish $k-$move, ie, a joint move that will assure that the max postload among the new links that are used by tasks of C is strictly less than the max preload among the links that were originally used by tasks of C. More precisely,

$$L \equiv \max_{j\in C}\left\{L_{\ell_j(t)}(t)\right\} > \max_{j\in C}\left\{L_{\ell_j(t+1)}(t+1)\right\} \equiv L',$$

where for any step r, $1 \leqslant r \leqslant T$, $\ell_j(r)$ is the link hosting task j and $L_\ell(r) = \sum_{j\in[n]:\ell_j(r)=\ell} w_j(\ell)$ is the load of link ℓ at r. By integrality of weights, $L' \leqslant L-1$.

For simplicity, let $\Lambda \equiv \{\ell \in [m] : \exists j \in C, \ell_j(t) = \ell\}$ be the set of *old* hosting links for tasks of C at step t, let $\ell^* \equiv \arg\max_{\ell\in\Lambda}\{L_\ell(t)\}$ be a link in Λ of maximum load (equal to L), and let $\Lambda' \equiv \{\ell \in [m] : \exists j \in C, \ell_j(t+1) = \ell\}$ be the set of *new* hosting links for tasks of C at step $t+1$. We consider an exponential potential $P(t) = \sum_{\ell\in[m]} D^{L_\ell(t)}$, where $D > 1$ is a positive constant to be specified later. A similar potential is used in [8] for unrelated parallel links and selfish 1-moves.

Using $L_{\ell^*}(t) = L$ and $L_\ell(t) \geqslant 0$ for the links in $\Lambda' \setminus \{\ell^*\}$, we obtain that

$$\sum_{\ell\in\Lambda\cup\Lambda'} D^{L_\ell(t)} \geqslant D^L + \sum_{\ell\in\Lambda\setminus(\Lambda'\cup\{\ell^*\})} D^{L_\ell(t)} + |\Lambda'\setminus\{\ell^*\}| \tag{1}$$

By selfishness of the k-move and integrality of weights, $\forall \ell \in \Lambda'$, $L_\ell(t+1) \leqslant L-1$. As for the remaining (old but not new) links, some load is removed from and no new load is added to them. Thus, $\forall \ell \in \Lambda\setminus\Lambda'$, $L_\ell(t+1) \leqslant L_\ell(t) - 1 \leqslant L-1$. Therefore,

$$\sum_{\ell\in\Lambda\cup\Lambda'} D^{L_\ell(t+1)} \leqslant |\Lambda'\cup\{\ell^*\}|\, D^{L-1} + \sum_{\ell\in\Lambda\setminus(\Lambda'\cup\{\ell^*\})} D^{L_\ell(t+1)} \tag{2}$$

Combining (1) and (2) and using that $\forall \ell \in \Lambda \setminus \Lambda'$, $L_\ell(t+1) < L_\ell(t)$ and that $\forall \ell \in [m] \setminus (\Lambda \cup \Lambda')$, $L_\ell(t+1) = L_\ell(t)$, we conclude that

$$\begin{aligned}P(t) - P(t+1) &\geqslant D^L + |\Lambda' \setminus \{\ell^*\}| - |\Lambda' \cup \{\ell^*\}|\, D^{L-1} \\ &= D^{L-1}(D - |\Lambda' \cup \{\ell^*\}|) + |\Lambda' \setminus \{\ell^*\}| \\ &\geqslant D - |\Lambda' \cup \{\ell^*\}| + |\Lambda' \setminus \{\ell^*\}| \\ &= D - 1\end{aligned}$$

provided that $D \geqslant |\Lambda' \cup \{\ell^*\}|$. Setting $D = k+1$ implies that $P(t) - P(t+1) \geqslant k$ and $P(t)$ is a generalized ordinal potential.

Initially, $P(0) \leqslant (k+1)^{W_{\text{tot}}}$ and at each step, the potential decreases by at least k. Since $P(t)$ remains positive, the maximum number of steps in any improvement path is less than $(k+1)^{W_{\text{tot}}}/k$. □

Remark. The above argument can be generalized to arbitrary (not necessarily integral) weights. The only difference is that $L - L'$ is bounded from below by some (sufficiently small) positive constant (instead of 1). The convergence rate however becomes much worse. The exponential potential essentially captures the power of the *lexicographic ordering* arguments for convergence.

Despite the existence of PNE, it is actually hard to compute one, even if we demand a very large number of coalitions of players, so long as we allow one coalition of large cardinality. This is shown in the next theorem:

Theorem 2. *For arbitrary static coalitional congestion games over identical parallel links among players with integer weights, it is* $\mathbf{NP}-$*complete to find a PNE, even if we enforce a number of coalitions* $k = 1 + n(1-\delta)$, *for any constant* $\delta \in (0, 1]$.

Proof. It is trivial to observe that for the case of two identical links and a single coalition $C_1 = \{w_1, \ldots, w_n\}$, any PNE is an optimum $2-$Partition of the weights $w_1, \ldots, w_n$, which is of course $\mathbf{NP}-$complete to construct.

Assume now that we demand any number $k = 1 + n(1-\delta)$ of coalitions for some constant $\delta \in (0, 1]$, where we assume wlog that $N \equiv n(1-\delta)$ is even. We shall construct a very simple instance that is hard to solve: This is shown again by reduction from an arbitrary $2-$Partition instance. In particular, fix the multiset of integer weights $A = \{w_i\}_{i\in[M]}$ of an arbitrary 2–Partition instance and let $W = \sum_{i\in[M]} w_i$, where $M \equiv \delta n$. We consider all these weights to be handled by a single coalition: $C_1 = A$. We then construct $k - 1 = N$ additional coalitions that are actually singletons. In particular, $\forall j \in [k]\setminus\{1\},\ C_j = \{1+W\}$. To complete the definition of the congestion game between these coalitions, consider a pair of identical parallel links to which the items of these coalitions must be assigned. Then it is easy to verify that any PNE of this coalitional congestion game must be such that each link possesses *exactly* $N/2$ of the large tasks: Clearly, if we had at least $N/2 + 1$ large weights to one link (say, the first link) then the other link would have at most $N/2 - 1$ large weights, and possibly also all the weights of A. Therefore, $L_1 \geqslant (N/2+1)(1+W) > (N/2+1)(1+W)-1 = N/2(1+W) - 1 + (1+W) \geqslant (N/2-1)(1+W) + W + (1+W) \geqslant L_2 + (1+W)$. So we conclude that in any PNE, each of the two links has exactly $N/2$ large weights. Now, the first coalition, C_1 must assign optimally (ie, minimizing the post load) its own M weights above the *equal loads* produced by the other coalitions and of course, if the eventual loads of the two links are equal then we have a YES instance of the initial $2-$Partition problem, otherwise we have a NO instance. □

Observe that (by Theorem 1) any improvement path of maximal length that allows arbitrary selfish $(\leqslant k)-$moves of players converges monotonically to a PNE that is also robust against any (even dynamically forming) coalition of at most k players. We call such a PNE a $k-$*robust* PNE. In the next section we explore better than the previous bounds on arbitrary improvement paths that end up with $k-$robust PNE, at least in the restricted case of identical machines.

4 Convergence Time to Robust Equilibria

In this section we initiate the discussion on the speed of convergence to PNE that are robust to arbitrary selfish k-moves (ie, selfish changes of coalitions that are of size at most k and can be formed even dynamically). Towards this direction, we seek for upper bounds on the convergence rate to a 2-robust PNE.

The rate of convergence assured by the exponential potential is rather poor. Our seek is for a possibly better rate of convergence, at least in the special case of identical parallel links. For identical parallel links and selfish 1-moves, the sum of the squares of the link loads, formally $F(t) = \sum_{\ell\in[m]}(L_\ell(t))^2$, is a weighted potential and guarantees pseudo-polynomial convergence to a PNE. On the other hand, for specific improvement strategies, one can use combinatorial arguments and show either linear (for the MAX WEIGHT PRIORITY–BEST RESPONSE strategy) or quadratic (for the FIFO PRIORITY–BEST RESPONSE strategy) convergence to an arbitrary PNE [8].

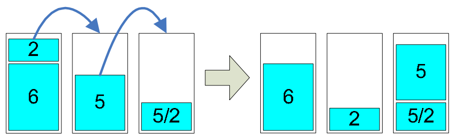

Fig. 1. A selfish 2-move increasing F by 1.

But we are also interested in the convergence rate of improvement paths with selfish 2-moves. Unfortunately, $F(t)$ is no longer a potential when we allow joint (selfish) moves of players, even if we consider only static coalitions of at most 2 players each. This is shown by the example of figure 1. Nevertheless, we shall consider two families of improvement paths with selfish $(\leqslant 2)$-moves, namely SMALLER COALITIONS FIRST and BEST OF THE BEST RESPONSE, for which $F(t)$ is a generalized potential.

4.1 Smaller Coalitions First

The first family consists of improvement paths that always give priority to selfish moves of smaller coalitions (therefore we shall call the strategy that creates these paths SMALLER COALITIONS FIRST strategy, $SCF(k)$ in short). In particular, we start by allowing (in arbitrary order) selfish $1-$moves, until no such move exists (ie, we are already at some $1-$robust PNE). Consequently we check whether there is a selfish $2-$move. If this is the case, then we make such a move, and then we perform again a maximal number of selfish $1-$moves. If not, we have already reached a $2-$robust PNE, and check whether there is a selfish $3-$move. If there are selfish $3-$moves, we allow one of them arbitrarily and consequently we perform (first) a maximal number of selfish $1-$moves and (from 1-robust PNE) selfish $2-$moves, before allowing another selfish $3-$move. We proceed in this manner, until we reach a $k-$robust PNE, in which case we stop.

In order to demonstrate the significance of this family of improvement paths, we shall show that indeed for $SCF(2)$ (which allows up to selfish $2-$moves and ends up with a $2-$robust PNE), $F(t)$ is indeed a generalized ordinal potential (ie, the difference $F(t) - F(t+1)$ is positive for selfish

($\leqslant 2$)$-$moves). This assures a pseudo-polynomial length for arbitrary improvement paths that mix selfish $1-$moves and $2-$moves arbitrarily, but always give priority to selfish $1-$moves.

Whenever a $2-$coalition of tasks (i, j) selfishly defects from their current hosts (ℓ_i, ℓ_j) respectively, towards two new links (ℓ_i', ℓ_j'), we denote this by $(i, j) \rhd (\ell_i, \ell_j) \mapsto (\ell_i', \ell_j')$. For $\ell_i \neq \ell_j$, the $2-$move $(i, j) \rhd (\ell_i, \ell_j) \mapsto (\ell_j, \ell_i)$ is called a $2-$*flip*, while the $2-$move $(i, j) \rhd (\ell_i, \ell_j) \mapsto (\ell_j, \ell_j')$ such that $\ell_j' \notin \{\ell_i, \ell_j\}$ is called a $2-$*chain*. We first explain why we may focus only on selfish $2-$flips and $2-$chains for the paths of $SCF(2)$:

Lemma 1. *The only possible selfish $2-$moves of any element of $SCF(2)$ are either $2-$flips or $2-$chains.*

Proof. Fix an arbitrary selfish $2-$move, while already being at a PNE (and therefore, there are no selfish $1-$moves). In particular, consider the selfish $2-$move $(i, j) \rhd (\ell_i, \ell_j) \mapsto (\ell_i', \ell_j')$.

If four distinct links are involved, then it is clear that this can be seen as two separate $1-$moves, and the $1-$move that involves the link of max preload (ie, before the $2-$move) is definitely a selfish $1-$move. But this cannot be true since we are already at a PNE.

Assume now that only two distinct links are involved. If the two players are initially hosted at the same link and then selfishly move to the same new link, this trivially implies the existence of selfish $1-$moves, which is not possible at PNEs.

Suppose finally that three distinct links are involved in a selfish $2-$move. It is easy to show that the cases of common destination link (ie, $(i, j) \rhd (\ell_i, \ell_j) \mapsto (\ell', \ell')$), and common original link (ie, $(i, j) \rhd (\ell, \ell) \mapsto (\ell_i', \ell_j')$) imply the existence of at least one selfish $1-$move at a PNE, a contradiction. For the case of common destination link, assume wlog that $L_{\ell_i}(t) \geqslant L_{\ell_j}(t)$. By selfishness and integrality we have: $L_{\ell_i}(t) \geqslant L_{\ell'}(t) + w_i + w_j + 1 > L_{\ell'}(t) + w_i + 1$. Similarly, for the case of common original link, we have: $L_{\ell}(t) \geqslant L_{\ell_i'}(t) + w_i + 1$.

It is very simple to construct PNE from which we can perform either a selfish $2-$flip, or a selfish $2-$chain. Therefore, we conclude that the only possible selfish $2-$moves that may appear in an improvement path of $SCF(2)$ are either $2-$flips or $2-$chains. □

The following lemma implies that if we restrict our attention to members of $SCF(2)$ that only allow selfish $2-$flips (and no 2–chains), then we are fair to the coalitional players, in the sense that no coalition can gain more by doing a $2-$chain rather than its corresponding $2-$flip. This is quite interesting, since a selfish $2-$flip seems much more "selfishly motivated" than a selfish $2-$chain: a selfish $2-$flip is essentially equivalent to (actually, slightly stronger than) the selfish $1-$move of the difference of two weights from the link with the heavier load to the link with the lesser load. In contrast, the $2-$chain cannot be decomposed into a sequence of selfish $1-$moves of weights, since one of the two transfers of weights may not be actually selfish.

Lemma 2. *At an arbitrary PNE, the existence of a selfish $2-$chain implies the existence of a selfish $2-$flip that assures at least the same improvement for the coalitional player that performs the move.*

Proof. Fix an arbitrary selfish $2-$chain from a PNE: $(i, j) \rhd (\ell_i, \ell_j) \mapsto (\ell_j, \ell_j')$. By the Nash property at the original configuration, $L_{\ell_i}(t) - w_i \leqslant L_{\ell_j'}(t)$. The coalitional cost *after* the $2-$chain is $\max\{L_{\ell_j}(t) + w_i - w_j, L_{\ell_j'}(t) + w_j\}$, while the coalitional cost after the $2-$flip $(i, j) \rhd (\ell_i, \ell_j) \mapsto (\ell_j, \ell_i)$ would be $\max\{L_{\ell_i}(t) - w_i + w_j, L_{\ell_j}(t) + w_i - w_j\}$. The latter cost is greater than the former only if $L_{\ell_i}(t) - w_i + w_j \geqslant L_{\ell_j'}(t) + w_j$, which implies that $L_{\ell_i}(t) - w_i > L_{\ell_j'}(t)$, a contradiction. Therefore, the $2-$flip is also selfish and the coalitional cost that it produces is at most equal to that of the selfish $2-$chain. □

The following theorem demonstrates the pseudo–polynomial convergence rate of any element of $SCF(2)$ that only allows selfish $2-$flips:

Theorem 3. *For an arbitrary system of identical parallel links and selfish users with integer weights, the function $F(t) = \sum_{\ell \in [m]} (L_\ell(t))^2$ is a generalized ordinal potential for all the members of $SCF(2)$ that only allow selfish $1-$moves and $2-$flips, and thus it assures their convergence to a $2-$robust PNE in at most $\frac{W_{\rm tot}{}^2}{2}$ steps.*

Proof. When a task i moves from link ℓ_i to ℓ'_i, $F(t)$ decreases by $2w_i[L_{\ell_i}(t) - (L_{\ell'_i}(t) + w_i)]$. Therefore, $F(t)$ is a weighted potential for selfish 1-moves (see also [8] for a generalization to uniformly related parallel links and [13] for a further generalization to weighted congestion games with linear delays).

So, let us assume that two tasks i and j are involved in a $2-$*flip*: $(i,j) \rhd (\ell_i, \ell_j) \mapsto (\ell_j, \ell_i)$. Wlog we suppose that $w_i \geqslant w_j$. Then,

$$\begin{aligned} F(t) - F(t+1) &= L^2_{\ell_i}(t) + L^2_{\ell_j}(t) - [L_{\ell_i}(t) - w_i + w_j]^2 - [L_{\ell_j}(t) + w_i - w_j]^2 \\ &= 2(w_i - w_j)\left[L_{\ell_i}(t) - (L_{\ell_j}(t) + w_i - w_j)\right] \end{aligned}$$

Hence, the 2-flip can be regarded as a 1-move of a weight $w_i - w_j$ from ℓ_i to ℓ_j, i.e. from the link hosting the heavier weight to the link hosting the lighter weight.

Initially, the cost of coalition (i,j) is $\max\{L_{\ell_i}(t), L_{\ell_j}(t)\}$. After the 2-flip, the coalition's cost is $\max\{L_{\ell_i}(t) - w_i + w_j, L_{\ell_j}(t) + w_i - w_j\}$. Assuming that the move is selfish, we obtain that

$$\max\{L_{\ell_i}(t), L_{\ell_j}(t)\} > \max\{L_{\ell_i}(t) - w_i + w_j, L_{\ell_j}(t) + w_i - w_j\} \tag{3}$$

Given that $w_i \geqslant w_j$, (3) holds only if $w_i > w_j$ and $L_{\ell_i}(t) > L_{\ell_j}(t)$. By integrality of weights, we obtain that $w_i - w_j \geqslant 1$ and $L_{\ell_i}(t) - L_{\ell_j}(t) \geqslant 1$. Therefore, $F(t) - F(t+1) \geqslant 4$ and $F(t)$ is a generalized ordinal potential for selfish 2-flips as well.

In overall, any improvement path of $SCF(2)$, which may only use a mixture of selfish $1-$moves and selfish $2-$flips, has for each such move a gain of at least 2. Since $F(0) \leqslant W_{\rm tot}{}^2$, we get the desired bound on the convergence rate. □

Unfortunately $F(t)$ is not a potential for selfish 2-chains (see e.g. Fig. 1). However, as shown in lemma 2, a selfish 2-chain at a 1-robust PNE implies the existence of a selfish 2-flip of no greater selfish cost. Therefore, we can always demand from the candidate coalition to perform a selfish 2-flip instead of a selfish 2-chain.

4.2 Best of the Best Response

Let (i,j) be a (possibly dynamic) coalition of two tasks. Wlog. we assume that $w_i \geqslant w_j$. Let ℓ_i and ℓ_j be the links to which i and j are currently assigned, and let $L_\ell^{-\{i,j\}}$ be the load of link ℓ in the current assignment due to the tasks in $[n] \setminus \{i,j\}$, i.e. tasks other than i and j.

The *best of the best response* of coalition (i,j) is to assign the heavier task i to the link minimizing the load caused by other tasks, and the lighter task j to the link minimizing the load caused by other tasks and i's new assignment[2]. Formally, task i is assigned to link $\ell'_i = \arg\min_{\ell \in [m]}\{L_\ell^{-\{i,j\}}\}$, and

[2] Any assignment of i and j that minimizes their maximum postload is a *best response* for coalition (i,j). Among all possible best responses, we focus on ones also minimizing the sum of i's and j's postload. Arguably, every coalition has a strong incentive to follow a *best of the best* response, i.e. a best response of minimum total postload.

task j is assigned to link

$$\ell'_j = \begin{cases} \arg\min_{\ell \in [m] \setminus \{\ell'_i\}} \{L_\ell^{-\{i,j\}}\} & \text{if } L_{\ell'_i}^{-\{i,j\}} + w_i \geqslant \min_{\ell \in [m] \setminus \{\ell'_i\}} \{L_\ell^{-\{i,j\}}\} \\ \ell'_i & \text{otherwise} \end{cases}$$

We call BEST OF THE BEST RESPONSE, $\mathrm{BoBR}(2)$ in short, the strategy that results in improvement paths where every coalition (i, j) switches to the links (ℓ'_i, ℓ'_j) defined as above. We prove that for dynamic coalitions of size (at most) 2, $F(t) = \sum_{\ell \in [m]} (L_\ell(t))^2$ is a generalized ordinal potential for $\mathrm{BoBR}(2)$. Assuming integer weights, this implies that $\mathrm{BoBR}(2)$ converges to a 2-robust PNE in no more than ${W_{\mathrm{tot}}}^2/2$ steps.

Theorem 4. *For every system of identical parallel links and selfish users with integer weights, the function $F(t) = \sum_{\ell \in [m]} (L_\ell(t))^2$ is a generalized ordinal potential for $\mathrm{BoBR}(2)$. Hence, $\mathrm{BoBR}(2)$ converges to a 2-robust PNE in at most $\frac{{W_{\mathrm{tot}}}^2}{2}$ steps.*

Proof. We show that $F(t)$ decreases by at least 2 when the selfish move $(i, j) \rhd (\ell_i, \ell_j) \mapsto (\ell'_i, \ell'_j)$ is performed. The intuition is that assigning (i, j) to (ℓ'_i, ℓ'_j) is both a coalition's best response and the assignment that minimizes the potential given that the remaining tasks do not move. Technically, the proof proceeds by case analysis.

If either $\ell_i = \ell'_i$ or $\ell_j = \ell'_j$, we have a selfish 1-move, for which $F(t)$ is a weighted potential and decreases by at least 2 (see also the proof of Theorem 3). Hence, we assume that $\ell_i \neq \ell'_i$ and $\ell_j \neq \ell'_j$ (i.e. both i and j move).

The most interesting case is when $\ell_i \neq \ell_j$ and $\ell'_i \neq \ell'_j$ (i.e. before and after the move, i and j are on different links). Then, ℓ'_j is the link with the second smallest load caused by tasks other than i and j. For simplicity, we let

$$x_i = L_{\ell_i}^{-\{i,j\}} = L_{\ell_i}(t) - w_i \qquad y_i = L_{\ell'_i}^{-\{i,j\}}$$
$$x_j = L_{\ell_j}^{-\{i,j\}} = L_{\ell_j}(t) - w_j \qquad y_j = L_{\ell'_j}^{-\{i,j\}}$$

Then, the change in the potential due to the selfish 2-move $(i, j) \rhd (\ell_i, \ell_j) \mapsto (\ell'_i, \ell'_j)$ is

$$\begin{aligned} F(t) - F(t+1) &= (x_i + w_i)^2 + (x_j + w_j)^2 + y_i^2 + y_j^2 - [x_i^2 + x_j^2 + (y_i + w_i)^2 + (y_j + w_j)^2] \\ &= 2w_i(x_i - y_i) + 2w_j(x_j - y_j) \end{aligned} \tag{4}$$

We emphasize that the first equality holds even if $\ell'_i = \ell_j$ (in which case $x_j = y_i$) or $\ell'_j = \ell_i$ (in which case $x_i = y_j$) or both $\ell'_i = \ell_j$ and $\ell'_j = \ell_i$ (in which case $x_i = y_j$ and $x_j = y_i$).

Since ℓ'_i minimizes the load caused by tasks other than i and j, $y_i \leqslant x_i$. If $y_j \leqslant x_j$, by selfishness of the 2-move, at least one of the inequalities $y_i \leqslant x_i$ and $y_j \leqslant x_j$ is strict the potential decreases. Then, $F(t) - F(t+1) \geqslant 2$ follows by integrality of weights.

On the other hand, if $y_j > x_j$, then $y_i = x_j < y_j \leqslant x_i$ because y_i is the smallest and y_j is the second smallest load caused by tasks other than i and j. Since the move is selfish, $x_i + w_i > y_j + w_j$. Then, by integrality of weights, at least one of $w_i \geqslant w_j + 1$ and $x_i \geqslant y_j + 1$ holds. In both cases, starting from (4) and using simple algebra, we conclude that $F(t) - F(t+1) \geqslant 2$.

The remaining cases are easily reduced to the special case of selfish 1-moves, for which $F(t)$ is a weighted potential (see also the proof of Lemma 1). More precisely,

- If i and j defect from the same link ℓ (i.e. $\ell_i = \ell_j = \ell$ and $\ell'_i \neq \ell'_j$), we regard w_i as moving first from ℓ to ℓ'_i and w_j as moving second from ℓ to ℓ'_j. At least the first move is selfish. The second move does not increase j's cost. Thus, the potential decreases by at least 2.

- The case that i and j move to the same link ℓ' (i.e. $\ell_i \neq \ell_j$ and $\ell'_i = \ell'_j = \ell'$) is similar. We regard w_i as moving first from ℓ_i to ℓ' and w_j as moving second from ℓ_j to ℓ'. As before, at least the first move is selfish and the potential decreases by at least 2.
- If i and j move from the same link ℓ to the same link ℓ' (i.e. $\ell_i = \ell_j = \ell$ and $\ell'_i = \ell'_j = \ell'$), we think of a single task of weight $w_i + w_j$ as moving selfishly from ℓ to ℓ'. The potential decreases by $2(w_i + w_j)(L_\ell(t) - L_{\ell'}(t+1)) \geqslant 2$.

In all cases, the potential decreases by at least 2 when coalition (i, j) switches to its BoBR. Initially $F(0) \leqslant {W_{\mathrm{tot}}}^2$. Since $F(t)$ remains positive, the length of any $\mathrm{BoBR}(2)$ improvement path is less than ${W_{\mathrm{tot}}}^2/2$. □

Remark. In the previous section, Lemma 2 shows that if the current assignment is a 1-robust PNE, a 2-flip is a BoBR for every coalition of size 2. Then, establishing that a selfish 2-flip decreases the potential is the crux of Theorem 3. In this section, we show that for every initial assignment, the potential decreases when a coalition of size 2 switches to its BoBR. If the current assignment is not a 1-robust PNE, a coalition's BoBR may be a 2-chain, which is nevertheless certain to decrease the potential.

4.3 The Generalized Fully Mixed Nash Equilibrium

Consequently we shall demonstrate the inexistence of Fully Mixed NE for coalitional congestion games with coalitions of cardinality at least 2. We shall also prove the existence of the so–called generalized fully mixed NE, for any coalitional congestion game on identical parallel links:

Lemma 3. *For the system of parallel identical links and users with arbitrary weights, there exists a fully mixed NE if and only if $n_j = 1$ for all the coalitions $j \in [k]$. If on the other hand there are coalitions of more than one users, then a generalized fully mixed NE that assures no conflict among players of the same coalition, always exists.*

Proof. If $n_j = 1$ for all $j \in [k]$ then the case reduces to the standard KP model [20], for which a fully mixed NE is guaranteed to exist.

Now fix a coalition $C_j, j \in [k]$, with $n_j > 1$. Suppose that there exists a fully mixed NE $\mathbf{p}$. Let $\theta_r(\mathbf{p_{-j}})$ be the expected load on link r caused by all coalitions except for C_j, with respect to $\mathbf{p}$. Let $\ell \in \arg\max_{r \in M} \theta_r(\mathbf{p_{-j}})$. Let σ_j be an optimum assignment of the loads of coalition j on the set of resources M, had j been the unique coalition. Since $n_j > 1$, it is clear that the contribution of C_j to the load of any link is at most $\mathsf{OPT}_j < W_j$. Let α_j be the assignment of the demands of coalition j where each w_i^j is assigned on resource ℓ, ie, $\alpha_j^i = \ell$ for all $i \in [n_j]$. Then clearly $\lambda_j(\mathbf{p}, \sigma_j) \leqslant \mathsf{OPT}_j + \theta_\ell(\mathbf{p_{-j}}) < W_j + \theta_r(\mathbf{p_{-j}}) = \lambda_j(\mathbf{p}, \alpha_j)$, a contradiction, because $\mathbf{p}$ is fully mixed and both assignments σ_j and α_j should give the same conditional expected selfish cost to coalition C_j.

If each coalition has at most m weights, we can easily determine a *generalized fully mixed* strategy, where the links for each coalition's weights are chosen uniformly at random but *without replacement* (ie, there is no chance for any two weights of the same player to be assigned to the same link). We call this profile the *generalized fully mixed NE (GFMNE).*

Now let $\mathbf{p}$ denote the GFMNE. To see that $\mathbf{p}$ is indeed a NE, observe that for any coalition $j \in [k]$, the expected load of any link $\ell \in [m]$ due to the other coalitions is $\theta_\ell(\mathbf{p_{-j}}) = \frac{W_{\mathrm{tot}} - W_j}{m}$ and the expected cost of player j is $\lambda_j(\mathbf{p}) = \frac{W_{\mathrm{tot}} - W_j}{m} + w_j^{\max}$. On the other hand, $\forall \alpha_j \in \Sigma_j$,

$$\lambda_j(\mathbf{p}, \alpha_j) = \max_{\ell \in S_j(\alpha_j)} \{\theta_\ell(\alpha_j) + \theta_\ell(\mathbf{p_{-j}})\} \geqslant w_j^{\max} + \frac{W_{\mathrm{tot}} - W_j}{m} = \lambda_j(\mathbf{p})$$

□

Remark. The construction of GFMNE can be easily extended to coalitions with more than m tasks. Every coalition with more than m tasks fixes an optimal m-partition of its tasks that is lexicographically minimum wrt the set load vector. Then every set of the m-partition is assigned to a link chosen uniformly at random without replacement. It is straightforward to verify that the corresponding profile is a generalized fully mixed NE.

5 Price of Anarchy in the Coalitional KP Model

Assume that there is a single coalition $C_1 = \widetilde{W}$, i.e. $k = 1$. Then, any NE is an optimum assignment of C_1 to M and vice versa. Hence in any NE σ, $\mathsf{SC}(\sigma) = \mathsf{OPT}$ and thus $\mathsf{R} = 1$. On the other hand, the case where $k = n$ reduces to the standard KP-model [20], for which $\mathsf{R} = \Theta\left(\frac{\log m}{\log \log m}\right)$ [19].

Then, we prove that for every $k \in [n]$, the price of anarchy is $\Theta(\min\{k, \frac{\log m}{\log\log m}\})$. The lower bound (Theorem 5) holds even for identical tasks and coalitions of equal cardinality. The upper bound (Theorem 6) holds for weighted tasks and arbitrary coalitions.

5.1 The Lower Bound

Theorem 5. *The price of anarchy is $\Omega(\min\{k, \frac{\log m}{\log\log m}\})$ even for identical tasks and coalitions of equal cardinality.*

Proof. We consider m identical parallel links and m unit size tasks partitioned into $k \geqslant 2$ coalitions each with $r \equiv m/k$ tasks (wlog we assume that m/k is an integer). We say that a coalition $j \in [k]$ *hits* a link $\ell \in [m]$ if j assigns at least one of its tasks to ℓ. We first prove a lower bound for the Coalitional Chains Model.

Lemma 4. *In the Coalitional Chains Model, when the number of coalitions is $k = m^{\varepsilon}$ for arbitrary constant $\varepsilon \in (0, 1]$, the price of anarchy is $\mathsf{R} = \Omega\left(\frac{\log m}{\log\log m}\right)$.*

Proof. It is easy to construct a very simple and natural (especially in scheduling literature) mixed profile that is indeed a NE and achieves the lower bound of $\Omega\left(\frac{\log m}{\log\log m}\right)$, for the case where the number of coalitions is $k = m^{\varepsilon}$, for arbitrary constant $\varepsilon \in (0, 1]$. This example is the following: Each coalition contains exactly $r = m/k = m^{1-\varepsilon}$ unit size tasks. We consider the following mixed profile $\mathbf{p}$ for the coalitions: Assuming that the m links form a cycle, each coalition $j \in [k]$ chooses uniformly at random a link $\ell_j \in [m]$ as its starting point, and then assigns its r weights to r consecutive links $\ell_j, \ell_j + 1(\bmod\ m), \ldots, \ell_j + r - 1(\bmod\ m)$.

We show first that this is indeed a NE: Observe that for any link $\ell \in [m]$, the expected load wrt any subset of coalitions using the profile $\mathbf{p}$ is the same: $\forall \ell \in [m], \theta_\ell(\mathbf{p}) = \sum_{j \in [k]} 1 \cdot \mathbb{P}\{j \text{ hits link } \ell\} = \sum_{j\in[k]} n^{-\varepsilon} = 1$ and $\forall \ell \in [m], \forall j \in [k], \theta_\ell(\mathbf{p_{-j}}) = \sum_{j' \in [k]\setminus\{j\}} 1 \cdot \mathbb{P}\{j' \text{ hits link } \ell\} = \sum_{j'\in[k]\setminus\{j\}} n^{-\varepsilon} = 1 - n^{-\varepsilon}$. Observe now that, since for any coalition $j \in [k]$ the expected load of the other players is exactly the same and there is no chance (wrt $\mathbf{p}$) that two balls of C_j fall into the same bin, the expected cost of C_j is $\lambda_j(\mathbf{p}) = 1 + 1 - n^{-\varepsilon} = 2 - n^{-\varepsilon}$. On the other hand, $\forall \sigma_j \in \Sigma_j, \lambda_j(\mathbf{p}, \sigma_j) = \max_{\ell \in S_j(\sigma_j)}\{\theta_\ell(\sigma_j) + \theta_\ell(\mathbf{p_{-j}})\} \geqslant 1 + 1 - n^{-\varepsilon} = \lambda_j(\mathbf{p})$, Therefore, we conclude that $\mathbf{p}$ is is a mixed NE for the coalitional game.

Next we prove that the coordination ratio for $\mathbf{p}$ is $\Omega\left(\frac{\log m}{\log\log m}\right)$. To see this, simply observe that the expected maximum load induced by $\mathbf{p}$ is lower bounded by the expected maximum load of *any* subset of links. We call a *chain* any set of r consecutive links in M (modulo m). Each player then chooses a chain independently and uniformly at random to allocate his r weights. Thus, a specific chain hits a link $\ell \in [m]$ if ℓ belongs to this chain. Consider the subset of links $M' \equiv \{1, r+1, 2r+1, \ldots, (k-1)r+1\}$. The links of this subset have the property that any chain will hit exactly one of the links in M'. Thus, the expected maximum load (wrt to the profile $\mathbf{p}$) among the links of M' equals the expected maximum load that we face when throwing k identical balls (the choices of the players' chains) to k identical bins (the links included in M'). But this is known that it equals $\Theta\left(\frac{\log k}{\log\log k}\right)$. Therefore, the social cost of $\mathbf{p}$ is lower bounded as follows: $\mathsf{SC}(\mathbf{p}) = \Omega\left(\frac{\log k}{\log\log k}\right) = \Omega\left(\frac{\log m^{\varepsilon}}{\log\log m^{\varepsilon}}\right) = \Omega\left(\frac{\varepsilon \log m}{\log \varepsilon + \log\log m}\right)$. The social optimum of this example is clearly 1. So, for any constant $1 \geqslant \varepsilon > 0$, we have that $\mathsf{R} = \Omega\left(\frac{\log m}{\log\log m}\right)$. □

General Case. As before, we consider m identical parallel links and m unit size tasks partitioned into $k \geqslant 2$ coalitions with $r \equiv m/k$ tasks each. The social optimum is equal to 1. The lower bound is established for the generalized fully mixed NE considered in Lemma 3.

In particular, let $\mathbf{p}$ be the mixed profile where every coalition chooses r links uniformly at random without replacement and assigns a task to every chosen link. In other words, every coalition assigns its tasks to any particular combination of r links with probability $1/\binom{m}{r}$. The probability that coalition j hits link ℓ is $1 - \binom{m-1}{r}/\binom{m}{r} = \frac{r}{m} = \frac{1}{k}$. Every coalition assigns at one of its tasks to every link. Lemma 3 shows that $\mathbf{p}$ is a NE for the coalitional game.

For every link $\ell \in [m]$ and every positive integer $\rho \leqslant k$, let $\mathcal{E}_\ell(\rho) \in \{0,1\}$ be the indicator variable of the event that link ℓ is hit by less than ρ coalitions in the outcome produced by $\mathbf{p}$. For convenience, we call a link ℓ *light* if it is hit by less than ρ coalitions, i.e. if $\mathcal{E}_\ell(\rho) = 1$, and *heavy* otherwise. The number of coalitions hitting a given link ℓ is the sum of k independent 0-1 random variables each becoming 1 with probability $1/k$. It is easy to calculate an upper bound on the probability that a given link is light.

$$\begin{aligned}
&\mathbb{P}\{\mathcal{E}_\ell(\rho) = 0\} \geqslant \binom{k}{\rho}\left(\frac{1}{k}\right)^{\rho}\left(1 - \frac{1}{k}\right)^{k-\rho} \geqslant \frac{1}{e\rho^{\rho}} \Rightarrow \\
&\mathbb{P}\{\mathcal{E}_\ell(\rho) = 1\} \leqslant 1 - \frac{1}{e\rho^{\rho}}
\end{aligned} \tag{5}$$

where we use that $\binom{k}{\rho} \geqslant (\frac{k}{\rho})^{\rho}$ and that for every integer $k \geqslant 2$, $(1 - \frac{1}{k})^{k-1} \geqslant \frac{1}{e}$.

The crucial observation is that the link lightness events are negatively associated[3]. Therefore, for every positive integer $\rho \leqslant k$,

$$\mathbb{P}\{\mathcal{E}_1(\rho) = \cdots = \mathcal{E}_m(\rho) = 1\} \leqslant \left(1 - \frac{1}{e\rho^{\rho}}\right)^{m} \leqslant e^{-\frac{m}{e\rho^{\rho}}} \tag{6}$$

The second inequality follows from $e^{-x} \geqslant 1 - x$. The first inequality is established in Lemma 5 below similarly to [17, Lemma 1].

Before we formally prove (6), we show that it indeed implies the lemma. In particular, we show that unless $\rho = O(\min\{k, \frac{\log m}{\log\log m}\})$, it is almost certain that in the configuration chosen wrt $\mathbf{p}$, at least one of the links is heavy. Indeed for every constant $c \geqslant 1$, there is a constant $\alpha \in (0,1)$ such that for all positive integers $\rho \leqslant \frac{\alpha \log m}{\log\log m}$, $\rho^{\rho} \leqslant \frac{m}{ce\log m}$ (note that α tends asymptotically to 1). Therefore, $e^{-\frac{m}{e\rho^{\rho}}} \leqslant m^{-c}$. In simple words, with probability at least $1 - m^{-c}$, there is some link hit

[3] The negative association of bin occupancies in the classical "balls and bins" experiment is established e.g. in [7].

by at least $\min\{k, \frac{\alpha \log m}{\log\log m}\}$ coalitions. If $k = \Omega(\frac{\log m}{\log\log m})$, the social cost of $\mathbf{p}$ is $\Theta(\frac{\log m}{\log\log m})$. If $k = o(\frac{\log m}{\log\log m})$, the social cost of $\mathbf{p}$ is $\Theta(k)$. Since the optimal social cost is 1, this implies the lower bound of Theorem 5.

The following lemma establishes that the link lightness events are negatively associated.

Lemma 5. *For any fixed positive integer $\rho \leqslant k$,*

$$\mathbb{P}\{\mathcal{E}_1(\rho) = \cdots = \mathcal{E}_m(\rho) = 1\} \leqslant \prod_{\ell=1}^{m} \mathbb{P}\{\mathcal{E}_\ell(\rho) = 1\}$$

Proof. For simplicity, we assume an arbitrary fixed threshold $\rho \leqslant k$ and denote the event that a link ℓ is light wrt ρ simply by $\mathcal{E}_\ell$ (instead of $\mathcal{E}_\ell(\rho)$). We first prove that for every $\ell \in \{2, \ldots, m\}$,

$$\mathbb{P}\{\mathcal{E}_1 = \cdots = \mathcal{E}_{\ell-1} = 1 | \mathcal{E}_\ell = 1\} \leqslant \mathbb{P}\{\mathcal{E}_1 = \cdots = \mathcal{E}_{\ell-1} = 1 | \mathcal{E}_\ell = 0\} \tag{7}$$

For all $l \in [m]$, let θ_l denote the number of coalitions hitting l in the outcome produced by $\mathbf{p}$. We fix an $\ell \in \{2, \ldots, m\}$. To establish (7), we consider the first ℓ links and fix the total number of hits in the links of $[\ell]$. In particular, let M be an arbitrary fixed integer such that $M = \sum_{l \in [\ell]} \theta_l$. We observe that

$$\begin{aligned}
&\mathbb{P}\left\{\mathcal{E}_1 = \cdots = \mathcal{E}_{\ell-1} = 1 | \mathcal{E}_\ell = 1;\ \textstyle\sum_{l \in [\ell]} \theta_l = M\right\} \\
&= \mathbb{P}\left\{\mathcal{E}_1 = \cdots = \mathcal{E}_{\ell-1} = 1 | \theta_\ell \leqslant \rho - 1;\ \textstyle\sum_{l \in [\ell]} \theta_l = M\right\} \\
&= \mathbb{P}\left\{\mathcal{E}_1 = \cdots = \mathcal{E}_{\ell-1} = 1 | \textstyle\sum_{l \in [\ell-1]} \theta_l \geqslant M - \rho + 1;\ \textstyle\sum_{l \in [\ell]} \theta_l = M\right\}
\end{aligned} \tag{8}$$

whereas

$$\begin{aligned}
&\mathbb{P}\left\{\mathcal{E}_1 = \cdots = \mathcal{E}_{\ell-1} = 1 | \mathcal{E}_\ell = 0;\ \textstyle\sum_{l \in [\ell]} \theta_l = M\right\} \\
&= \mathbb{P}\left\{\mathcal{E}_1 = \cdots = \mathcal{E}_{\ell-1} = 1 | \theta_\ell \geqslant \rho;\ \textstyle\sum_{l \in [\ell]} \theta_l = M\right\} \\
&= \mathbb{P}\left\{\mathcal{E}_1 = \cdots = \mathcal{E}_{\ell-1} = 1 | \textstyle\sum_{l \in [\ell-1]} \theta_l \leqslant M - \rho;\ \textstyle\sum_{l \in [\ell]} \theta_l = M\right\}
\end{aligned} \tag{9}$$

Now it is clear that the probability given in (8) is upper bounded by the one given in (9), since in both cases we assume the same amount of hits in the links of $[\ell]$, but for (8) we assume that more hits (than those assumed by (9)) regard links of $[\ell - 1]$. This implies (7) because it holds for any total number of hits in the links of $[\ell]$, and thus for any vector of coalitional hits in $[\ell]$.

Using (7), we obtain that for every $\ell \in \{2, \ldots, m\}$,

$$\begin{aligned}
\mathbb{P}\{\mathcal{E}_1 = \cdots = \mathcal{E}_{\ell-1} = 1\} &= \mathbb{P}\{\mathcal{E}_1 = \cdots = \mathcal{E}_{\ell-1} = 1 | \mathcal{E}_\ell = 1\} \cdot \mathbb{P}\{\mathcal{E}_\ell = 1\} \\
&\quad + \mathbb{P}\{\mathcal{E}_1 = \cdots = \mathcal{E}_{\ell-1} = 1 | \mathcal{E}_\ell = 0\} \cdot \mathbb{P}\{\mathcal{E}_\ell = 0\} \\
&\geqslant \mathbb{P}\{\mathcal{E}_1 = \cdots = \mathcal{E}_{\ell-1} = 1 | \mathcal{E}_\ell = 1\}
\end{aligned} \tag{10}$$

Therefore,

$$\begin{aligned}
&\mathbb{P}\{\mathcal{E}_1 = \cdots = \mathcal{E}_m = 1\} = \mathbb{P}\{\mathcal{E}_1 = \cdots = \mathcal{E}_{m-1} = 1 | \mathcal{E}_m = 1\} \cdot \mathbb{P}\{\mathcal{E}_m = 1\} \\
&\leqslant \mathbb{P}\{\mathcal{E}_1 = \cdots = \mathcal{E}_{m-1} = 1\} \cdot \mathbb{P}\{\mathcal{E}_m = 1\} \leqslant \cdots \leqslant \prod_{\ell=1}^{m} \mathbb{P}\{\mathcal{E}_\ell = 1\}
\end{aligned}$$

where the inequalities are due to a recursive application of (10). □

Lemma 5 establishes the negative association of the link lightness events and concludes the proof of the lower bound. □

5.2 The Upper Bound

In this section, we establish an asymptotically tight upper bound on the Price of Anarchy for the coalitional game on identical parallel links. The approach is similar to the approach of [12, Theorem 7], which bounds the price of anarchy in the non-cooperative KP-model with identical users and related parallel links.

Theorem 6. *For every NE* $\mathbf{p}$*,* $\mathsf{SC}(\mathbf{p}) \leqslant O(\min\{k, \frac{\log m}{\log\log m}\})\mathsf{OPT}$.

Proof. We start with some notation and an easy fact about NE. Let $S_j(\mathbf{p_j}) = \bigcup_{\sigma_j : p_j(\sigma_j)>0} S_j(\sigma_j)$ denote the set of links with positive probability of getting some task of coalition j in mixed strategy $\mathbf{p_j}$. Let $\theta_\ell^{\max}(\mathbf{p_j}) = \max_{\sigma_j : p_j(\sigma_j)>0}\{\theta_\ell(\sigma_j)\}$ be the maximum load assigned to link ℓ by coalition j in mixed strategy $\mathbf{p_j}$.

Proposition 1. *Let* $\mathbf{p}$ *be a NE. Then,*
(a) For every coalition j *and every link* $\ell \in [m]$*,* $\theta_\ell^{\max}(\mathbf{p_j}) \leqslant \lambda_j(\mathbf{p})$.
(b) For every coalition j *and every link* $\ell \in S_j(\mathbf{p_j})$*,* $\theta_\ell(\mathbf{p}) \leqslant \lambda_j(\mathbf{p})$.

Proof. Let σ_j be the pure strategy of coalition j with $p_j(\sigma_j) > 0$ that maximizes $\theta_\ell(\sigma_j)$. Hence, $\theta_\ell(\sigma_j) = \theta_\ell^{\max}(\mathbf{p_j})$ and either $\theta_\ell^{\max}(\mathbf{p_j}) = 0$ or $\ell \in S_j(\sigma_j)$. We recall that $\lambda_j(\mathbf{p}, \sigma_j) = \lambda_j(\mathbf{p})$ by Definition 2 (Nash Equilibrium).

If $\theta_\ell^{\max}(\mathbf{p_j}) = 0$, (a) follows from $\lambda_j(\mathbf{p}) \geqslant 0$. If $\ell \in S_j(\sigma_j)$, (a) holds because

$$\theta_\ell^{\max}(\mathbf{p_j}) = \theta_\ell(\sigma_j) \leqslant \theta_\ell(\sigma_j) + \theta_\ell(\mathbf{p_{-j}}) \leqslant \lambda_j(\mathbf{p}, \sigma_j) = \lambda_j(\mathbf{p})$$

As for (b), let $\alpha_j \in \Sigma_j$ be any pure strategy with $p_j(\alpha_j) > 0$. By the definition of σ_j, $\theta_\ell(\alpha_j) \leqslant \theta_\ell(\sigma_j)$ and

$$\theta_\ell(\alpha_j) + \theta_\ell(\mathbf{p_{-j}}) \leqslant \theta_\ell(\sigma_j) + \theta_\ell(\mathbf{p_{-j}}) \leqslant \lambda_j(\mathbf{p}) \tag{11}$$

Therefore,

$$\theta_\ell(\mathbf{p}) = \sum_{\alpha_j \in \Sigma_j} p_j(\alpha_j)(\theta_\ell(\alpha_j) + \theta_\ell(\mathbf{p_{-j}})) \leqslant \sum_{\alpha_j : p_j(\alpha_j)>0} p_j(\alpha_j)\lambda_j(\mathbf{p}) = \lambda_j(\mathbf{p})$$

The inequality follows by applying (11) for every pure strategy α_j with $p_j(\alpha_j) > 0$. □

An Upper Bound on Selfish Costs. We proceed by bounding the selfish cost of any coalition in a NE in terms of the optimal cost.

Lemma 6. *Let* $\mathbf{p}$ *be any NE. For every coalition* j*,* $\lambda_j(\mathbf{p}) \leqslant (2 - \frac{1}{m})\mathsf{OPT}$.

Proof. For sake of contradiction, let us assume that there is a coalition j with $\lambda_j(\mathbf{p}) > (2 - \frac{1}{m})\mathsf{OPT}$. Let $\sigma_j \in \Sigma_j$ be any pure strategy with $p_j(\sigma_j) > 0$. By Definition 2 (Nash Equilibrium), σ_j is a best response to $\mathbf{p_{-j}}$ and $\lambda_j(\mathbf{p}, \sigma_j) = \lambda_j(\mathbf{p})$. Also let

$$M_h(\mathbf{p}, \sigma_j) = \{\ell \in S_j(\sigma_j) : \theta_\ell(\sigma_j) + \theta_\ell(\mathbf{p_{-j}}) = \lambda_j(\mathbf{p}, \sigma_j)\}\,, \text{ and let}$$
$$M_l(\mathbf{p}, \sigma_j) = \{\ell \in [m] : \theta_\ell(\sigma_j) + \theta_\ell(\mathbf{p_{-j}}) \leqslant (1 - \tfrac{1}{m})\mathsf{OPT}\}\,.$$

For convenience, let $m_h = |M_h(\mathbf{p}, \sigma_j)|$, and let $m_l = |M_l(\mathbf{p}, \sigma_j)|$.

Claim. If $\lambda_j(\mathbf{p}, \sigma_j) > (2 - \frac{1}{m})\mathsf{OPT}$, then $m_h \leqslant m_l$.

Before we prove the claim, we show that it indeed implies the lemma. If $\lambda_j(\mathbf{p}, \sigma_j) > (2-\frac{1}{m})\mathsf{OPT}$ and $m_h \leqslant m_l$, we can reassign a task of coalition j from every link in $M_h(\mathbf{p}, \sigma_j)$ to a different link in $M_l(\mathbf{p}, \sigma_j)$. The new strategy σ'_j obtained by these reassignments has $\lambda_j(\mathbf{p}, \sigma'_j) < \lambda_j(\mathbf{p}, \sigma_j)$. More specifically, some load is removed from every link $\ell \in M_h(\mathbf{p}, \sigma_j)$ and thus $\theta_\ell(\sigma'_j) < \theta_\ell(\sigma_j)$. On the other hand, the load of every link $\ell \in M_l(\mathbf{p}, \sigma_j)$ increases by at most OPT and $\theta_\ell(\sigma'_j) + \theta_\ell(\mathbf{p_{-j}}) \leqslant (2-\frac{1}{m})\mathsf{OPT}$. Therefore, σ_j is not a best response to $\mathbf{p_{-j}}$, a contradiction to the hypothesis that $\mathbf{p}$ is a NE.
Proof of the claim. The proof is by contradiction. Let us assume that $\lambda_j(\mathbf{p}, \sigma_j) > (2-\frac{1}{m})\mathsf{OPT}$ and that $m_h > m_l$. Since both m_h, m_l are integers, $m_h \geqslant m_l + 1$. Then,

$$\begin{aligned}
W_{\text{tot}} &= \sum_{\ell \in [m]} (\theta_\ell(\sigma_j) + \theta_\ell(\mathbf{p_{-j}})) \\
&> [(2-\tfrac{1}{m})m_h + (1-\tfrac{1}{m})(m - m_h - m_l)]\mathsf{OPT} \\
&= [m_h + m - 1 - (1-\tfrac{1}{m})m_l]\mathsf{OPT} \\
&\geqslant m\mathsf{OPT} \geqslant W_{\text{tot}} \,,
\end{aligned}$$

a contradiction. The first inequality holds because every link in $\ell \in M_h(\mathbf{p}, \sigma_j)$ has $\theta_\ell(\sigma_j)+\theta_\ell(\mathbf{p_{-j}}) > (2-\frac{1}{m})\mathsf{OPT}$ and every link $\ell \notin M_h(\mathbf{p}, \sigma_j) \cup M_l(\mathbf{p}, \sigma_j)$ has $\theta_\ell(\sigma_j) + \theta_\ell(\mathbf{p_{-j}}) > (1-\frac{1}{m})\mathsf{OPT}$. The second inequality holds because $m_h \geqslant m_l + 1$ implies that $m_h - 1 - (1-\frac{1}{m})m_l \geqslant 0$. The last inequality follows from $\mathsf{OPT} \geqslant W_{\text{tot}}/m$. □

Remark. Lemma 6 is essentially tight even for 2 coalitions. For a lower bound, we consider an instance of m identical links, $m(m-1)$ small tasks of unit weight, and 2 large tasks of weight m. The tasks are partitioned into 2 coalitions, the first comprising the small tasks and the second comprising the large tasks. The optimal solution has cost $m+1$. A PNE of social cost $2m$ is formed by assigning the large tasks to a single link and m small tasks to each of the remaining links.

Bounding the Price of Anarchy. Then, we show that for every NE $\mathbf{p}$, $\mathsf{SC}(\mathbf{p}) \leqslant O(k)\mathsf{OPT}$ (cf. Lemma 7) and $\mathsf{SC}(\mathbf{p}) \leqslant O(\frac{\log m}{\log\log m})\mathsf{OPT}$ (Lemma 8).

Lemma 7. *For every NE* $\mathbf{p}$*,* $\mathsf{SC}(\mathbf{p}) \leqslant (2-\frac{1}{m})\, k\, \mathsf{OPT}$.

Proof. Let $\sigma = (\sigma_j)_{j \in [k]}$ be any configuration with positive probability in $\mathbf{p}$. For every link ℓ,

$$\theta_\ell(\sigma) = \sum_{j=1}^{k} \theta_\ell(\sigma_j) \leqslant \sum_{j=1}^{k} \theta_\ell^{\max}(\mathbf{p_j}) \leqslant \sum_{j=1}^{k} \lambda_j(\mathbf{p}) \leqslant k\,(2-\tfrac{1}{m})\mathsf{OPT}$$

The first inequality follows from the definition of $\theta_\ell^{\max}$. Then, we use that $\theta_\ell^{\max}(\mathbf{p_j}) \leqslant \lambda_j(\mathbf{p}) \leqslant (2-\frac{1}{m})\mathsf{OPT}$ by Proposition 1.a and Lemma 6. □

Lemma 8. *For every NE* $\mathbf{p}$*,* $\mathsf{SC}(\mathbf{p}) \leqslant O(\frac{\log m}{\log\log m})\mathsf{OPT}$.

Proof. For every link ℓ and every coalition j, let X_ℓ^j be the random variable describing the actual load assigned to ℓ by j. The random variable X_ℓ^j is equal to $\theta_\ell(\sigma_j)$ if j adopts pure strategy σ_j, which happens with probability $p_j(\sigma_j)$. Consequently, $\mathbb{E}\left\{X_\ell^j\right\} = \theta_\ell(\mathbf{p_j})$ and X_ℓ^j takes values in $[0, \theta_\ell^{\max}(\mathbf{p_j})] \subseteq [0, (2-\frac{1}{m})\mathsf{OPT}]$ (recall that $\theta_\ell^{\max}(\mathbf{p_j}) \leqslant (2-\frac{1}{m})\mathsf{OPT}$ by Proposition 1.a and

Lemma 6). For a given link ℓ, the random variables $(X_\ell^j)_{j\in[k]}$ are independent because the coalitions select their strategies independently.

For every link ℓ, let $X_\ell = \sum_{j=1}^k X_\ell^j$ be the random variable describing the actual load assigned to link ℓ. X_ℓ is the sum of k independent random variables with values in $[0, (2-\frac{1}{m})\mathsf{OPT}]$. By linearity of expectation, $\mathbb{E}\{X_\ell\} = \theta_\ell(\mathbf{p})$. By Proposition 1.b and Lemma 6, $\mathbb{E}\{X_\ell\} \leqslant (2-\frac{1}{m})\mathsf{OPT}$.

Applying the standard Hoeffding bound[4] with $w = (2-\frac{1}{m})\mathsf{OPT}$ and $t = e\rho\,(2-\frac{1}{m})\mathsf{OPT}$, we obtain that for every link ℓ and every $\rho \geqslant 1$,

$$\mathbb{P}\left\{X_\ell > e\rho\,(2-\tfrac{1}{m})\mathsf{OPT}\right\} \leqslant \rho^{-e\rho}$$

We apply the union bound and conclude that for every $\rho \geqslant 1$,

$$\mathbb{P}\left\{\max_{\ell\in[m]}\{X_\ell\} > e\rho\,(2-\tfrac{1}{m})\mathsf{OPT}\right\} \leqslant m\rho^{-e\rho} \tag{12}$$

The social cost of $\mathbf{p}$ is equal to the expected maximum load assigned to some link. Formally,

$$\mathsf{SC}(\mathbf{p}) = \mathbb{E}\left\{\max_{\ell\in[m]}\{X_\ell\}\right\} = \sum_{i=0}^{\infty} \mathbb{P}\left\{\max_{\ell\in[m]}\{X_\ell\} > i\right\} \tag{13}$$

The last equality holds because the random variable $\max_{\ell\in[m]}\{X_\ell\}$ takes only non-negative integer values. It is well-known that for every non-negative integer-valued random variable X, $\mathbb{E}\{X\} = \sum_{i=0}^{\infty}\mathbb{P}\{X > i\}$ (see e.g. [24, Proposition C.7]).

Using (12) and (13), we obtain that for every $\rho_0 \geqslant 2$,

$$\begin{aligned}
\mathsf{SC}(\mathbf{p}) &\leqslant e\rho_0\,(2-\tfrac{1}{m})\mathsf{OPT} + \sum_{i=e\rho_0\,(2-\frac{1}{m})\mathsf{OPT}}^{\infty} \mathbb{P}\left\{\max_{\ell\in[m]}\{X_\ell\} > i\right\} \\
&\leqslant e\rho_0\,(2-\tfrac{1}{m})\mathsf{OPT} + \sum_{\rho=\rho_0}^{\infty} e\,(2-\tfrac{1}{m})\mathsf{OPT}\,\mathbb{P}\left\{\max_{\ell\in[m]}\{X_\ell\} > e\rho\,(2-\tfrac{1}{m})\mathsf{OPT}\right\} \\
&\leqslant e\,(2-\tfrac{1}{m})\mathsf{OPT}\left(\rho_0 + \sum_{\rho=\rho_0}^{\infty} m\rho^{-e\rho}\right) \\
&\leqslant e\,(2-\tfrac{1}{m})\mathsf{OPT}\left(\rho_0 + 2m\rho_0^{-e\rho_0}\right)
\end{aligned}$$

We first use (13) and $\mathbb{P}\left\{\max_{\ell\in[m]}\{X_\ell\} > i\right\} \leqslant 1$ for $i = 0, \dots, e\rho_0\,(2-\frac{1}{m})\mathsf{OPT} - 1$. Then, we change the variable of the summation. The third inequality follows from (12). For the last inequality, we first observe that $\rho^{-e\rho} \leqslant \rho_0^{-e\rho}$ for all $\rho \geqslant \rho_0$. Hence,

$$\sum_{\rho=\rho_0}^{\infty} m\rho^{-e\rho} \leqslant m\sum_{\rho=\rho_0}^{\infty}\left(\rho_0^{-e}\right)^{\rho} \leqslant 2m\rho_0^{-e\rho_0}$$

because $\rho_0^{-e} \leqslant 1/2$. For $\rho_0 = \frac{2\log m}{\log\log m}$, we obtain that $\rho_0^{-e\rho_0} \leqslant m^{-1}$ for every $m \geqslant 4$. Therefore, $\mathsf{SC}(\mathbf{p}) \leqslant 4e\,(\frac{\log m}{\log\log m} + 1)\mathsf{OPT}$. □

The upper bound of Theorem 6 follows immediately by combining the upper bounds of Lemma 7 and Lemma 8. □

[4] We use the standard version of Hoeffding bound [16]: 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, $\forall t > 0, \mathbb{P}\{X > t\} \leqslant (\frac{e\,\mathbb{E}\{X\}}{t})^{t/w}$.

6 Pure Nash Equilbria in Coalitional Congestion Games

In this section, we assume an arbitrary network topology and use the weighted total delay of the coalition as our notion of selfish cost. We first prove that a weighted congestion game with coalitions is an exact potential game if the edge delays are linear.

Theorem 7. *Every weighted network congestion game with coalitions and linear edge delays is an exact potential game.*

Proof. Let $G(V, E)$ be a directed network with a linear delay $d_e(x) = a_e x + b_e$, $a_e, b_e \geqslant 0$, associated with each edge $e \in E$, let $\{C_1, \dots, C_k\}$ be a set of coalitions, and let $\sigma = (\sigma_j^i)_{j\in[k]}$ be any configuration. The potential of σ is $\Phi(\sigma) = \frac{1}{2}[C(\sigma) + W(\sigma)]$, where

$$C(\sigma) = \sum_{e\in E} \theta_e(\sigma) d_e(\theta_e(\sigma)) = \sum_{e\in E} (a_e\theta_e^2(\sigma) + b_e\theta_e(\sigma)) \text{ , and}$$

$$W(\sigma) = \sum_{e\in E}\sum_{j=1}^{k} \theta_e(\sigma_j) d_e(\theta_e(\sigma_j)) = \sum_{e\in E}\sum_{j=1}^{k} (a_e\theta_e^2(\sigma_j) + b_e\theta_e(\sigma_j))$$

Let j be a coalition that switches its strategy from σ_j to σ'_j, and let $\sigma' = \sigma_{-j} \oplus \sigma'_j$ be resulting configuration. The difference in the total delay of coalition j is[5]:

$$\begin{aligned}
\tau_j(\sigma') - \tau_j(\sigma) &= \sum_{e\in E} \theta_e(\sigma'_j)[a_e(\theta_e(\sigma_{-j}) + \theta_e(\sigma'_j)) + b_e] - \sum_{e\in E} \theta_e(\sigma_j)[a_e(\theta_e(\sigma_{-j}) + \theta_e(\sigma_j)) + b_e] \\
&= \sum_{e\in E} (\theta_e(\sigma'_j) - \theta_e(\sigma_j))[a_e(\theta_e(\sigma_{-j}) + \theta_e(\sigma'_j) + \theta_e(\sigma_j)) + b_e] \qquad (14)
\end{aligned}$$

The difference in the potential function components is:

$$\begin{aligned}
C(\sigma') - C(\sigma) &= \sum_{e\in E} [a_e(\theta_e(\sigma_{-j}) + \theta_e(\sigma'_j))^2 + b_e(\theta_e(\sigma_{-j}) + \theta_e(\sigma'_j))] \\
&\quad - \sum_{e\in E} [a_e(\theta_e(\sigma_{-j}) + \theta_e(\sigma_j))^2 + b_e(\theta_e(\sigma_{-j}) + \theta_e(\sigma_j))] \\
&= \sum_{e\in E} (\theta_e(\sigma'_j) - \theta_e(\sigma_j))[a_e(2\,\theta_e(\sigma_{-j}) + \theta_e(\sigma'_j) + \theta_e(\sigma_j)) + b_e] \qquad (15)
\end{aligned}$$

and

$$\begin{aligned}
W(\sigma') - W(\sigma) &= \sum_{e\in E} [a_e\theta_e^2(\sigma'_j) + b_e\theta_e(\sigma'_j)] - \sum_{e\in E} [a_e\theta_e^2(\sigma_j) + b_e\theta_e(\sigma_j)] \\
&= \sum_{e\in E} (\theta_e(\sigma'_j) - \theta_e(\sigma_j))[a_e(\theta_e(\sigma'_j) + \theta_e(\sigma_j)) + b_e] \qquad (16)
\end{aligned}$$

Combing (15), (16), and (14), we obtain that

$$\begin{aligned}
[C(\sigma') + W(\sigma')] - [C(\sigma) + W(\sigma)] &= \sum_{e\in E} 2\,(\theta_e(\sigma'_j) - \theta_e(\sigma_j))[a_e(\theta_e(\sigma_{-j}) + \theta_e(\sigma'_j) + \theta_e(\sigma_j)) + b_e] \\
&= 2\,(\tau_j(\sigma') - \tau_j(\sigma))
\end{aligned}$$

Therefore, $\Phi(\sigma') - \Phi(\sigma) = \tau_j(\sigma') - \tau_j(\sigma)$. Hence, Φ is an exact potential for weighted congestion games with coalitions and linear delays. □

[5] In the following, we repeatedly use that for every pure assignment α, coalition j, and edge e, $\theta_e(\alpha) = \theta_e(\alpha_{-j}) + \theta_e(\alpha_j)$.

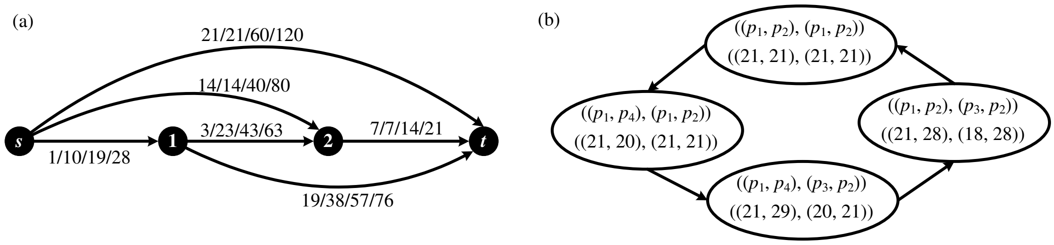

Fig. 2. (a). (a). A coalitional network congestion game that does not have the Finite Improvement Property. Consider two coalitions both consisting of 2 identical users routing their traffic from s to t. There are 4 different $s-t$ paths: $p_1 = (s,t)$, $p_2 = (s,2,t)$, $p_3 = (s,1,2,t)$, and $p_4 = (s,1,t)$. The notation $a/b/c/d$ means that a single user has delay a, two users have delay b, three users have delay c, and four users have delay d. A PNE for this game is $((p_1,p_2),(p_1,p_3))$. (b). An improvement cycle of length 4 demonstrating that this game does not have the Finite Improvement Property. Every node corresponds to a pure assignment. The delays of the paths used by each coalition are shown under the assignment. Every arrow corresponds to a defection improving the delay of the coalition changing its strategy.

Remark. Theorem 7 can be easily generalized to weighted (not necessarily network) congestion games with linear resource delays. Theorem 7 can be regarded as a generalization of [13, Theorem 1] to coalitions. In [13, Theorem 1], the same potential function is shown to be a weighted potential function for weighted congestion games with linear resource delays.

On the other hand, we give a simple example of a unweighted single-commodity network congestion game with coalitions of equal cardinality and 2-wise linear edge delays that does not have the *Finite Improvement Property* (FIP). By [23, Lemma 2.5], this game does not admit a potential function of any kind.

Lemma 9. *There exist instances of unweighted single-commodity network congestion games with coalitions of equal cardinality and edge delays either linear of 2-wise linear which do not have the Finite Improvement Property.*

Proof. The Nash dynamics (ie, the directed graph with pure assignments as nodes and coalition-delay-improving defections as edges) for the game of Fig. 2.a contains a cycle of length 4. Therefore, the game does not have the Finite Improvement Property.

Suppose that both coalitions start with (p_1,p_2). The delay along both p_1 and p_2 is 21 and the total delay of both coalitions is 42. The first coalition improves its total delay to 41 by switching to (p_1,p_4). The delay of the second coalition remains 41. Then, the second coalition switches to (p_3,p_2) and improves its delay from 42 to 41. This increases the delay of the first coalition from 41 to 50. Then, the first coalition switches back to (p_1,p_2) and improves its delay to 49. This increases the delay of the second coalition from 41 to 46. Finally, the second coalition switches back to (p_1,p_2) and improves its delay to 42 (see also Fig. 2.b). $\square$

References

1. Awerbuch B., Azar Y., Epstein A. The price of routing unsplittable flow. In *Proc. of the 37th ACM Symp. on Th. of Comp. (STOC '05)*, pages 57–66, 2005.
2. G. Christodoulou and E. Koutsoupias. The Price of Anarchy of Finite Congestion Games. In *Proc. of the 37th ACM Symp. on Th. of Comp. (STOC '05)*, pages 67–73, 2005.

3. Cominetti R., Correa J.R., Stier Moses N.E. Network games with atomic players. In *Proc. of the 33rd Int. Col. on Automata, Lang. and Progr. (ICALP '06)*. Springer-Verlag, 2006. (to appear).
4. Correa J.R., Schulz A.S., Stier Moses N.E. Computational complexity, fairness, and the price of anarchy of the maximum latency problem. In *Proc. of the 10th Conf. on Int. Progr. and Comb. Opt. (IPCO '04) LNCS 3064*, pages 59–73, 2004.
5. Correa J.R., Schulz A.S., Stier Moses N.E. On the inefficiency of equilibria in congestion games. In *Proc. of the 11th Conf. on Int. Progr. and Comb. Opt. (IPCO '05)*, pages 167–181, 2005.
6. Czumaj A., Vöcking B. Tight bounds for worst-case equilibria. In *Proc. of the 13th ACM-SIAM Symp. on Discr. Alg. (SODA '02)*, pages 413–420, 2002.
7. D. Dubhashi and D. Ranjan. Balls and Bins: A Study in Negative Dependence. Technical Report RS-96-25, BRICS, 1996.
8. Even-Dar E., Kesselman A., Mansour Y. Convergence time to nash equilibria. In *Proc. of the 30th Int. Col. on Automata, Lang. and Progr. (ICALP '03)*, pages 502–513. Springer-Verlag, 2003.
9. Even-Dar E., Mansour Y. Fast convergence of selfish rerouting. In *Proc. of the 16th ACM-SIAM Symp. on Discr. Alg. (SODA '05)*, pages 772–781. SIAM, 2005.
10. Fabrikant A., Papadimitriou C., Talwar K. The complexity of pure nash equilibria. In *Proc. of the 36th ACM Symp. on Th. of Comp. (STOC '04)*, 2004.
11. 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 Int. Col. on Automata, Lang. and Progr. (ICALP '03)*, pages 514–526. Springer-Verlag, 2003.
12. 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 Int. Col. on Automata, Lang. and Progr. (ICALP '02)*, pages 123–134. Springer-Verlag, 2002.
13. Fotakis D., Kontogiannis S., Spirakis P. Selfish unsplittable flows. *Theoretical Computer Science*, 348((2-3)):226–239, December 2005. Special Issue dedicated to ICALP 2004 (TRACK-A).
14. Fotakis D., Kontogiannis S., Spirakis P. Symmetry in network congestion games: Pure equilibria and anarchy cost. In *Proc. of the 3rd W. on Approx. and Online Algorithms (WAOA '05)*, pages 161–175. Springer-Verlag, 2006.
15. Hayrapetyan A., Tardos E., Wexler T. The effect of collusion in congestion games. In *Proc. of the 38th ACM Symp. on Th. of Comp. (STOC '06)*, 2006. (to appear).
16. Hoeffding W. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
17. A. Kamath, R. Motwani, K. Palem, and P. Spirakis. Tail Bounds for Occupancy and the Satisfiability Threshold Conjecture. *Random Structures and Algorithms*, 7(1):59–80, 1995.
18. Kontogiannis S., Spirakis P. Atomic selfish routing in networks: A survey. In *Proc. of the 1st W. on Internet and Net. Econ. (WINE '05)*, LNCS 3828, pages 989–1002. Springer, 2005.
19. Koutsoupias E., Mavronicolas M., Spirakis P. Approximate equilibria and ball fusion. *Theory of Computing Systems*, 36(6):683–693, 2003. Special Issue devoted to SIROCCO'02.
20. Koutsoupias E., Papadimitriou C. Worst-case equilibria. In *Proc. of the 16th Annual Symp. on Theor. Aspects of Comp. Sci. (STACS '99)*, pages 404–413. Springer-Verlag, 1999.
21. Libman L., Orda A. Atomic resource sharing in noncooperative networks. *Telecommunication Systems*, 17(4):385–409, 2001.
22. Mavronicolas M., Spirakis P. The price of selfish routing. In *Proc. of the 33rd ACM Symp. on Th. of Comp. (STOC '01)*, pages 510–519, 2001.
23. Monderer D., Shapley L. Potential games. *Games & Econ. Behavior*, 14:124–143, 1996.
24. Motwani R., Raghavan P. *Randomized Algorithms*. Cambridge Univ. Press, 1995.
25. Nash J. F. Noncooperative games. *Annals of Mathematics*, 54:289–295, 1951.
26. Orda A., Rom R., Shimkin N. Competitive routing in multi-user communication networks. *IEEE/ACM Trans. on Networking*, 1(5):510–521, 1993.
27. Roughgarden T. Selfish routing with atomic players. In *Proc. of the 16th ACM-SIAM Symp. on Discr. Alg. (SODA '05)*, pages 973–974, 2005.