Theoretical Computer Science 974 (2023) 114094

Contents lists available at ScienceDirect

Theoretical Computer Science

journal homepage: www.elsevier.com/locate/tcs

Congestion games with priority-based scheduling ☆,☆☆

Vittorio Bilò *, Cosimo Vinci

Department of Mathematics and Physics "Ennio De Giorgi", University of Salento, Provinciale Lecce-Arnesano, P.O. Box 193, 73100 Lecce, Italy

A R T I C L E   I N F O

Article history:
Received 19 October 2022
Received in revised form 16 July 2023
Accepted 23 July 2023
Available online 31 July 2023
Communicated by M. Mavronicolas

Keywords:
Pure Nash equilibria
Price of anarchy
Price of stability

A B S T R A C T

We reconsider atomic and non-atomic affine congestion games under the assumption that players are partitioned into p priority classes and resources schedule their users according to a priority-based policy, breaking ties uniformly at random. We derive tight bounds on both the price of anarchy and the price of stability as a function of p, revealing an interesting separation between the general case of $p \geq 2$ and the priority-free scenario of $p = 1$. In fact, while in absence of priorities the worst-case prices of anarchy and stability of non-atomic games are lower than their counterparts in atomic ones, the two classes share the same bounds when $p \geq 2$. Moreover, while the worst-case price of stability is lower than the worst-case price of anarchy in atomic games with no priorities, their values become equal when $p \geq 2$. Said differently, the presence of priorities simultaneously irons out any combinatorial difference between atomic and non-atomic requests and among different pure Nash equilibria to produce a unique representative worst-case situation. Notably, our results keep holding even under singleton strategies. Besides being of independent interest, priority-based scheduling shares tight connections with online load balancing and finds a natural application within the theory of coordination mechanisms and cost-sharing policies for congestion games. Under this perspective, a number of possible research directions also arise.

1. Introduction

Measuring the inefficiency caused by selfish behavior in non-cooperative systems is one of the leading research directions in Algorithmic Game Theory [46]. Selfishness dictates that system users (from now on, *players*) wish to minimize their usage cost only. This personal objective usually stands in contrast with that of optimizing the system performance, expressed by some global function such as the sum of the players' costs or the maximum player cost. With this respect, the *price of anarchy* [43], comparing the worst-case equilibrium solution with the global optimum, and the *price of stability* [2], focusing on the best-case instead, have been introduced as inefficiency measures.

In this work, we are going to consider scenarios which can be illustrated through the following example. Imagine you are spending a holiday with some colleagues in a hotel by the beach. The hotel entertainment department wishes to offer complimentary activities for the group on the last day of your stay. You can choose between horse-riding and a fishing experience. As there is only one available horse and one available boat, which can accommodate one passenger only, sched-

☆ This article belongs to Section A: Algorithms, automata, complexity and games, Edited by Paul Spirakis.

☆☆ A preliminary version of this paper appeared in the Proceedings of 13th International Symposium on Algorithmic Game Theory [12].

* Corresponding author.
E-mail addresses: vittorio.bilo@unisalento.it (V. Bilò), cosimo.vinci@unisalento.it (C. Vinci).

https://doi.org/10.1016/j.tcs.2023.114094

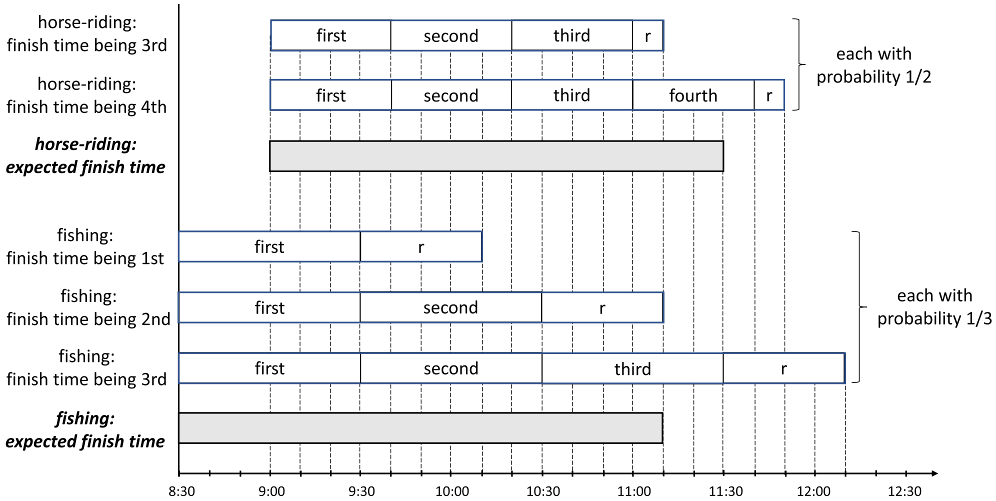

Fig. 1. A pictorial representation of the possible different schedules, and of the expected resulting ones, arising in our motivating example. Letter "r" stands for (time to) return (to the hotel).

ules on the two resources must be arranged to fulfill all the requests (your group is small enough, so that everyone can be satisfied within the day, even if you all were to choose the same activity). The schedules will be determined by giving priority to the owners of membership cards, breaking ties by lotteries. Both activities are equally interesting to you, so your desire is to finish as soon as possible, in order to have plenty of time for a relaxing bath before the return trip. Horse-riding starts at 9 a.m. and takes 40 minutes plus 10 minutes to return from the stables; the fishing experience starts at 8:30 a.m. and takes 1 hour plus 40 minutes to return from the dock. Your group amounts to six people: two of them are gold-card-owners and the others are silver-card-owners. You own a silver card, the two gold-card-owners and one silver-card-owner are going horse-riding, *which activity will you choose?*

If you go for horse-riding, you are equally likely to be third or fourth in the schedule, so your expected finish time will be $\frac{1}{2}(11{:}10) + \frac{1}{2}(11{:}50) = 11{:}30$. If you choose the fishing experience, you are equally likely to be first, second or third in the schedule, so your expected finish time will be $\frac{1}{3}(10{:}10) + \frac{1}{3}(11{:}10) + \frac{1}{3}(12{:}10) = 11{:}10$. So, you should choose the fishing experience (see Fig. 1 for a pictorial representation of the different possible schedules). If we assume that all other players have the same objective as you, i.e., finishing as soon as possible under indifference between the two options, a *pure Nash equilibrium* [45], that is a solution in which no player improves by changing her choice, is realized. In fact, the gold-card-owners have an expected finish time of $\frac{1}{2}(9{:}50) + \frac{1}{2}(10{:}30) = 10{:}10$. If one of them switches to the fishing experience, she still finishes at 10:10. The unique silver-card-owner choosing horse-riding finishes at 11:10. By switching to the fishing experience, her expected finish time grows to $\frac{1}{4}(10{:}10) + \frac{1}{4}(11{:}10) + \frac{1}{4}(12{:}10) + \frac{1}{4}(13{:}10) = 11{:}40$. All players choosing the fishing experience are in the same situation as you, and we have already checked that this option is the best possible one, given the choices of the others. So, nobody improves by deviating.

It is worth observing that the presence of priority-based scheduling tremendously impacts on your final decision. In fact, if priorities were not used, then the following happens. If you go for horse-riding, you are equally likely to be first, second, third or fourth in the schedule, so your expected finish time will be $\frac{1}{4}(9{:}50) + \frac{1}{4}(10{:}30) + \frac{1}{4}(11{:}10) + \frac{1}{4}(11{:}50) = 10{:}50$. If you choose the fishing experience, nothing changes, so your expected finish time remains 11:10. So, in this case, you end up choosing horse-riding. This immediately implies that the pure Nash equilibrium considered above loses its stability property in the priority-free scenario, i.e., the set of pure Nash equilibria of a game may be influenced by the use of priority scheduling. On the other hand, however, the presence or absence of priorities does not change the set of finish times of all players. In fact, for any outcome in which, let's say, 3 players go for horse-riding, there is always one player ending at 9:50, another ending at 10:30 and the last one ending at 11:10. The presence of priorities, indeed, only changes the order of the schedule and, in turn, the personal cost of the players. Thus, although the set of equilibria may be influenced by the use of priority scheduling, the socially optimal solution remains the same in both models. This allows for a fair comparison among the inefficiencies caused by different scheduling policies, and poses the intriguing question of determining to what extent the presence of priorities impacts on the price of anarchy and the price of stability of a given game.[1]

The above example falls within the general class of *affine congestion games*. Congestion games, introduced in [51], model scenarios in which a finite set of players compete for the usage of a finite set of resources, and the cost that every player

[1] This property holds as, in our model, the finish times of all players have the same weight in the social optimum. We will not consider scenarios with weighted players in this work.

pays for using a resource only depends on the amount of its users (a.k.a. the resource congestion).[2] In affine congestion games, the resource cost functions, called *latency functions*, are affine in their congestion. There are two fundamental models of congestion games, namely *atomic* and *non-atomic games*, which differ on the way in which players are interpreted. In atomic games, the congestion that every player causes on a resource is non-negligible and normalized to one (as in our illustrating example), whereas, in the non-atomic variant, each player is responsible for an infinitesimally small congestion effect (consider for example a single car in a six-lane long road during rush hours). Roughly speaking, a non-atomic game can be seen as an atomic game in which meta-players are allowed to arbitrarily split their requests along different sets of resources. Atomic congestion games always admit pure Nash equilibria for any type of latency functions [51], and the same holds also for non-atomic games under the assumption that the latency functions are continuous and non-decreasing [54]. A fundamental additional feature of non-atomic games is that all pure Nash equilibria share the same global cost, so that the price of anarchy and the price of stability collapse into a unifying representative.

Priority-free affine atomic congestion games, that is, affine congestion games under *randomized scheduling*, have been considered in [50] and [13]. In particular, for the global function of minimizing the sum of the players' costs, [50] shows that the price of anarchy is $5/3$, while the price of stability is bounded to $1+1/\sqrt{5}\approx 1.447$ in [13]. For non-atomic games, a simple adaptation from [54] yields a tight bound of $4/3$ on both the price of anarchy and the price of stability. To the best of our knowledge, no particular results are known under the assumption of priority-based scheduling, if one excludes the case in which there exists a fixed (possibly even resource-specific) ordering of the players and each resource processes the players' requests according to this ordering (*ordered-based scheduling*). This is essentially the same as assuming that there are as many priority classes as players, with each player belonging to a different priority class, so that there are no ties to be resolved, and the final schedule becomes deterministic. For this extreme case, results given by [27] and [31] imply a price of anarchy of 4 for both atomic and non-atomic games, respectively.

We try to fill this gap by studying (affine) congestion games with priority-based scheduling. In these games, we assume that n players are partitioned into p priority classes, with $1 \le p \le n$, and that, on every resource, all players of priority class c are scheduled before any player of class $c' > c$ (the lower the class, the higher the priority), while players of the same class are scheduled in a random order. Hence, the cost that player i experiences on resource r is a function of two parameters: *(i)* the position that i occupies in the schedule of r, and *(ii)* the latency function of r, i.e., how fast r processes its requests. More precisely, denoted by $\ell_r(x)$ the latency function of r, i suffers a cost equal to $\ell_r(k)$ when occupying the kth position in the schedule of r. Parameter *(i)*, which depends on i's priority class, is a random variable. Thus, the cost of i becomes the sum of the expected cost she experiences on every selected resource. Observe that the two previously considered scenarios, i.e., randomized scheduling and ordered-based scheduling, are re-obtained as the extreme cases of $p=1$ and $p=n$, respectively. Our aim is thus that of determining how the price of anarchy and the price of stability vary as a function of p.

1.1. Our results

When adding priority-based scheduling in a congestion game, it is no longer true that the cost that each player pays on a resource depends only on the number of its users. Rather, it starts depending on two parameters: the number of users being in the same priority class and the number of users with higher priorities. It follows that existential guarantees of pure Nash equilibria cannot be inherited from the priority-free scenario and need to be reproved. Thus, as our first result, we show that both atomic and non-atomic priority-based affine congestion games admit pure Nash equilibria. As for non-atomic games it also turns out that all pure Nash equilibria share the same social cost, it follows that the price of anarchy and the price of stability coincide within this class. For the sake of simplicity, all proofs are specifically tailored to deal with affine latencies, but they can be trivially generalized to arbitrary functions obeying mild assumptions. The techniques we exploit here are extensions of state-of-the-art ideas from the literature, such as the definition of potential functions for atomic games [44,51] and the first order condition for optimality in convex optimization for non-atomic ones [54].

Having shown the existence of pure Nash equilibria in both models, our main result is then the derivation of tight bounds for both the price of anarchy and the price of stability of priority-based affine congestion games as a function of p and with respect to the sum of the players' costs. These bounds, which are reported in Table 1, are tight even under singleton strategies and reveal an interesting separation between the general case of $p \ge 2$ and the priority-free scenario of $p=1$. In fact, while in absence of priorities the worst-case prices of anarchy and stability of non-atomic games are lower than their counterparts in atomic ones (prices of anarchy and stability of $4/3$ [54] vs. a price of anarchy of $5/3$ [50] and a price of stability of $1+1/\sqrt{5}$ [13], respectively), the two classes share the same bounds when $p \ge 2$. Moreover, while the worst-case price of stability is lower than the worst-case price of anarchy in atomic games with no priorities ($1+1/\sqrt{5}$ vs. $5/3$), their values become equal when $p \ge 2$. Said differently, the presence of priorities simultaneously irons out any

[2] A famous and well-studied generalization of congestion games, called *weighted congestion games*, assumes that players have different weights and so the cost that every player pays for using a resource becomes dependent on the total weight of its users. In such a setting, the traditional model of [51] gets also called *unweighted congestion games* and is equivalent to weighted games in which all weights are equal and normalized to one. In this work, we adopt the classical nomenclature in Game Theory which indicates with the general name of *congestion games* the original model of unweighted congestion games introduced in [51].

Table 1
The price of anarchy (PoA) of priority-based affine congestion games for some values of p. For $p = 1$, the bound 4/3 [54] holds for non-atomic games, while the bound 5/3 [50] holds for atomic ones. For $p \geq 2$, the bound is unique for both types of games and holds even for the price of stability (PoS) under singleton strategies.

p	PoA	p	PoA	p	PoA	p	PoA
1	$\{4/3, 5/3\}$	6	2.9683	11	3.4576	20	3.7625
2	2	7	3.1063	12	3.5137	30	3.8756
3	2.3248	8	3.2196	13	3.5617	40	3.9239
4	2.5875	9	3.3133	14	3.603	50	3.9487
5	2.7984	10	3.3916	15	3.6389	∞	4

combinatorial difference between atomic and non-atomic requests and among pure Nash equilibria, to produce a unique representative worst-case situation.[3] To the best of our knowledge, this is the first example of such a unified behavior.

The technique we exploit to derive the upper bounds is the primal-dual method introduced in [6] and based on pairs of primal/dual formulations. The application of this framework, however, requires a clever analysis of the set of dual constraints (a countable set for atomic games, an uncountable set for non-atomic ones), to extract a representative subset of worst-case constraints whose satisfaction yields a non-linear program defining the desired price of anarchy bound. By numerically solving this program, which has $p + 1$ constraints, one obtains the values reported in Table 1. Then, by leveraging the complementary slackness conditions, the existence of load balancing games, i.e. games in which players can only choose among single resources, attaining matching lower bounds, which can be even extended to the price of stability, is derived. We observe that the construction of the lower bounding instances is parametrized on an optimal solution for the non-linear program which need not to be known in order to show the mere existence of a matching lower bound.

An interesting application of our results falls within the problem of online scheduling with related machines and identical jobs. Assume we have an online scheduling problem P, where we are given a set of m related machines, with machine i having speed s_i, and an input sequence of n unit-length jobs (with n not known in advance), each coming with an associated set of machines where it can be processed. Suppose also that the input sequence is divided into p subsequences of jobs ($1 \leq p \leq n$ does not need to be known) and that, for each $1 \leq c \leq p$, when the c-th subsequence arrives, the sets of allowable machines of all jobs in the subsequence are immediately revealed (thus, the traditional setting in which jobs arrive one at time coincides with the case of $p = n$). Now interpret P as an atomic affine congestion game with p priority classes, where each subsequence is seen as a priority class and each machine of speed s is a resource having latency function equal to $\ell(x) = x/s$. As the cost of a job belonging to class c is not influenced by jobs of higher classes, it is easy to see (and a formal evidence is provided in the proof of Theorem 1) that one can inductively construct a pure Nash equilibrium for the congestion game yielded by the jobs of class c upon a pure Nash equilibrium for the game induced by all jobs belonging to classes smaller than c, so as to obtain a pure Nash equilibrium for the whole game P. As a pure Nash equilibrium for load balancing congestion games can be computed in polynomial time [38]), our results provide a polynomial time algorithm for online scheduling with related machines and identical jobs arranged into p sub-sequences and characterize its competitive ratio as a function of p.

1.2. Related work

The majority of the literature devoted to congestion games (e.g. [3,23,43,51,54] and subsequent work) assumes that all users experience the same cost on a same resource. This can be interpreted as the outcome of a *round-robin* scheduling policy under the assumption that requests are processed according to a time-sharing policy organized in such a way that all requests are completed (almost) simultaneously. Note that this requires preemption of the requests. Alternatively, this cost model can be thought as if requests are released only after all of them have been processed. The study of the efficiency of (pure) Nash equilibria in congestion games under the round-robin scheduling policy initiated with the seminal papers [2,3,23,43,54]. Since then, many results have been obtained in the literature under different generalizations or specializations, see [1,5,7,6,9,17,20,21,53,55]. When the resources have affine latency functions, the price of anarchy is 5/2 [23] and the price of stability is $1 + 1/\sqrt{3} \approx 1.577$ [17,22] for atomic games; for non-atomic ones both metrics are equal to 4/3 [54].

Other approaches, for which preemption is not necessary, consider Smith's Rule [27], the first-in first-out policy [31] and the random policy [13,27,43,50]. Smith's Rule and the random policy, in particular, can be seen as the specialization of priority-based scheduling obtained when $p = n$ and $p = 1$, respectively. More generally, Farzad et al. [31] focus on a broad family of non-preemptive scheduling policies, which may be even resource-specific, but always producing a total ordering of the users of every resource. For affine latencies, they show a price of anarchy of 4 for non-atomic games and a price of anarchy of 17/3 for atomic ones. Observe that, while the price of anarchy of non-atomic games coincides with the one we

[3] We stress that, by using the price of anarchy and the price of stability as measures of efficiency, our comparison between atomic and non-atomic games is based on worst-case instances. This does not rule out the possibility, and indeed there are examples where this happens, that the performance of pure Nash equilibria of a particular non-atomic game may improve when transforming it to an atomic one.

derive in this case when $p = n$, this is not the case for atomic games, where the price of anarchy in our model is much lower. This is due to the fact that, in [31], a different cost function is considered. In fact, while we assume that a player scheduled at position k on a resource with latency function $\ell(x)$ pays a cost of $\ell(k)$, [31] assumes a cost of $\int_{k-1}^{k} \ell(x)dx$. While this diversity is inconsequential in non-atomic games, for atomic ones a different cost model, with different efficiency bounds, arises. Finally, [31] also considers generalizations to polynomial latency functions and to weighted players.

In general, much better bounds are possible when either preemption or randomization is allowed. In fact, for atomic games, a preemptive scheduling policy yielding a price of anarchy of 5/2 is derived in [27], while, under the random policy, the price of anarchy drops to 5/3 [50] and the price of stability to $1 + 1/\sqrt{5} \approx 1.447$ [13]. Thus, as in this setting the efficiency of pure Nash equilibria is tremendously influenced by the chosen strategy, the study of the efficiency of different scheduling policies, besides being interesting per se, plays a fundamental role also in the theory of coordination mechanisms and cost-sharing policies for congestion games.

A *coordination mechanism* [24] is a local policy rule that each resource applies to schedule its assigned requests, while a cost-sharing policy [30,34,35] is a rule determining how the cost of a resource has to be shared among its users. Both machineries are usually used with the aim of mitigating the inefficiencies caused by selfish behavior. Coordination mechanisms for congestion games have been considered in [4,15,16,18,24,26,39]. Other techniques developed to cope with selfish behavior in congestion games are taxation, studied in [10,19,28,33,36,37,40,47–49,56], and Stackelberg strategies, considered in [11,29,32,41,52,56].

Tight connections between (singleton) affine congestion games and greedy algorithms for (online) scheduling problems have been noted and investigated in several papers, such as [8,9,14,17,25,27,42,55,57]. The mostly related to our work is [27], where a selfish scheduling game with unrelated machines and weighted jobs is considered. In this setting, the cost of a job is defined as its weight times its completion time. The authors consider different scheduling policies, such as Smith's Rule, a pre-emptive policy based on proportional time-sharing, and a non-uniform random policy. The price of anarchy of these three strategies is shown to be equal to 4, 2.618 (which drops to 2.5 for uniform weights) and 2.134, respectively. Then, some of the game-theoretical properties of these policies are suitably exploited to design a $(2 + \epsilon)$-approximation algorithm for the centralized problem of minimizing the sum of the players' costs. This last result has been improved to 1.81 in [18].

1.3. Paper organization

The next section introduces the model and all preliminary concepts and definitions. The technical part of the paper is divided into two sections. The first (Section 3) contains the existential results of pure Nash equilibria, while the second (Section 4) presents the characterization of their efficiency, which is the main contribution of this work. Finally, we conclude in Section 5 by discussing possible future research directions.

2. Model and definitions

For an integer $k \geq 1$, denote by $[k] := \{1, \dots, k\}$ the set of the first k positive integers. Moreover, set $[0] := \emptyset$.

2.1. Atomic games

For any integer $p \geq 1$, a *priority-based affine atomic congestion game with p priority classes* $\Gamma_p^a = ([n], R, (\mathcal{S}_i)_{i\in[n]}, (\alpha_r, \beta_r)_{r\in R}, (P_c)_{c\in[p]})$, where superscript a stands for *atomic*, is defined by a finite set $[n]$ of $n \geq 2$ players, a finite set R of resources, a strategy set $\mathcal{S}_i \subseteq 2^R \setminus \emptyset$ for each player $i \in [n]$, two coefficients $\alpha_r \geq 0$ and $\beta_r \geq 0$ for each resource $r \in R$ and a priority class $P_c \subseteq [n]$ for each $c \in [p]$ such that $\cup_{c\in[p]} P_c = [n]$ and $P_c \cap P_{c'} = \emptyset$ for each $c, c' \in [p]$ with $c \neq c'$, i.e., the sets $P_1, \dots, P_p$ realize a partition of $[n]$. We use $c(i)$ to refer to the priority class of player i, i.e., $c(i) = j$ if and only if $i \in P_j$.

Denote by $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)$ the strategy profile in which each player $i \in [n]$ chooses strategy $\sigma_i \in \mathcal{S}_i$. For a strategy profile $\boldsymbol{\sigma}$, a priority class $c \in [p]$ and a resource $r \in R$, let $n_r^c(\boldsymbol{\sigma}) = |\{i \in P_c : r \in \sigma_i\}|$ be the number of players belonging to class c selecting resource r in $\boldsymbol{\sigma}$, $n_r^{<c}(\boldsymbol{\sigma}) = \sum_{c'\in[c-1]} n_r^{c'}(\boldsymbol{\sigma})$ be the number of players belonging to any class $c' < c$ selecting resource r in $\boldsymbol{\sigma}$ and $n_r(\boldsymbol{\sigma}) = \sum_{c\in[p]} n_r^c(\boldsymbol{\sigma})$ be the congestion of resource r in $\boldsymbol{\sigma}$, i.e., the number of its users.

The cost that a player experiences on resource r when she occupies the kth position in the schedule of r (say i is the kth user of r) is equal to $\alpha_r k + \beta_r$ (affine latency functions). Thus, the expected cost of player i in $\boldsymbol{\sigma}$ is defined as

$$
\begin{aligned}
cost_i(\boldsymbol{\sigma}) &= \sum_{r\in\sigma_i} \sum_{k\in[n_r(\boldsymbol{\sigma})]} ((\alpha_r k + \beta_r) \cdot \Pr[i \text{ is the } k\text{th user of } r]) \\
&= \sum_{r\in\sigma_i} \left(\alpha_r \left(n_r^{<c(i)}(\boldsymbol{\sigma}) + \frac{1}{n_r^{c(i)}(\boldsymbol{\sigma})} \sum_{k\in[n_r^{c(i)}(\boldsymbol{\sigma})]} k \right) + \beta_r \right) \\
&= \sum_{r\in\sigma_i} \left(\alpha_r \left(n_r^{<c(i)}(\boldsymbol{\sigma}) + \frac{n_r^{c(i)}(\boldsymbol{\sigma}) + 1}{2} \right) + \beta_r \right).
\end{aligned}
$$

The utilitarian social cost, from now on simply the *social cost*, of $\boldsymbol{\sigma}$ is defined as the sum of the expected cost of all players in $\boldsymbol{\sigma}$, thus equal to

$$\begin{aligned}\mathsf{SC}(\boldsymbol{\sigma}) &= \sum_{i\in[n]} cost_i(\boldsymbol{\sigma})\\ &= \sum_{r\in R}\sum_{k\in[n_r(\boldsymbol{\sigma})]}(\alpha_r k+\beta_r)\\ &= \sum_{r\in R}\left(\alpha_r\frac{n_r(\boldsymbol{\sigma})(n_r(\boldsymbol{\sigma})+1)}{2}+\beta_r n_r(\boldsymbol{\sigma})\right),\end{aligned}$$

where the last equality easily follows by observing that, for each $r\in R$ with $n_r(\boldsymbol{\sigma})$ users, there is exactly one player occupying the kth position in the schedule of r for each $k\in[n_r(\boldsymbol{\sigma})]$. We shall denote by $\boldsymbol{\sigma}^*=(\sigma_1^*,\ldots,\sigma_n^*)$ the *social optimum* of Γ_p^a, that is, the strategy profile minimizing the social cost.

We shall focus on the notion of pure Nash equilibrium which is defined as follows.

Definition 1. A strategy profile $\boldsymbol{\sigma}$ is a pure Nash equilibrium for Γ_p^a if, for each $i\in[n]$ and $S\in\mathcal{S}_i$, $cost_i(\boldsymbol{\sigma})\le cost_i(\boldsymbol{\sigma}_{-i},S)$.[4]

By the above definition, given a pure Nash equilibrium $\boldsymbol{\sigma}$ and a social optimum $\boldsymbol{\sigma}^*$ for Γ_p^a, the following inequality holds for each $i\in[n]$:

$$\begin{aligned}&\sum_{r\in\sigma_i}\left(\alpha_r\left(n_r^{<c(i)}(\boldsymbol{\sigma})+\frac{n_r^{c(i)}(\boldsymbol{\sigma})+1}{2}\right)+\beta_r\right)\\ &\quad-\sum_{r\in\sigma_i^*}\left(\alpha_r\left(n_r^{<c(i)}(\boldsymbol{\sigma})+\frac{n_r^{c(i)}(\boldsymbol{\sigma})+2}{2}\right)+\beta_r\right)\le 0.\end{aligned}\tag{1}$$

2.2. Non-atomic games

For any integer $p\ge 1$, a *priority-based affine non-atomic congestion game with p priority classes* $\Gamma_p^{na}=([n],R,(f_i)_{i\in[n]},(\mathcal{S}_i)_{i\in[n]},(\alpha_r,\beta_r)_{r\in R},(P_c)_{c\in[p]})$, where the superscript *na* stands for *non-atomic*, has the same definition of its atomic counterpart with a different interpretation on the set of players and on how they handle their requests. For every $i\in[n]$, in fact, there is an amount of flow f_i belonging to priority class $c(i)$ that needs to be assigned to strategies in $\mathcal{S}_i$ in an arbitrarily splittable way. Thus, informally speaking, every flow can be interpreted as a set of infinitely many players, all belonging to the same priority class, each contributing a negligible amount to the congestion of the used resources. Let $m_i=|\mathcal{S}_i|$ denote the number of strategies available to the ith flow and set $\mathcal{S}_i:=\{S_{i,1},\ldots,S_{i,m_i}\}$. In this setting, a strategy profile is identified by a tuple $\boldsymbol{\sigma}=(\sigma_{1,1},\ldots\sigma_{1,m_1},\ldots,\sigma_{n,1},\ldots\sigma_{n,m_n})$, where, for every $i\in[n]$ and $j\in[m_i]$, $\sigma_{i,j}\ge 0$ denotes the fraction of the ith flow assigned to $S_{i,j}$. We shall only consider feasible strategy profiles, i.e., strategy profiles such that $\sum_{j\in[m_i]}\sigma_{i,j}=f_i$ for each $i\in[n]$. We overload the notation of $\boldsymbol{\sigma}$ for the sake of analyzing both atomic and non-atomic games under the same framework. To this aim, we also denote by $n_r^c(\boldsymbol{\sigma})=\sum_{i\in P_c}\sum_{j\in[m_i]:r\in S_{i,j}}\sigma_{i,j}$ the total amount of flow of priority class c assigned to resource r in $\boldsymbol{\sigma}$. Similarly, we define $n_r^{<c}(\boldsymbol{\sigma})=\sum_{c'\in[c-1]}n_r^{c'}(\boldsymbol{\sigma})$ and $n_r(\boldsymbol{\sigma})=\sum_{c\in[p]}n_r^c(\boldsymbol{\sigma})$. The difference is that, while all these quantities are non-negative reals here, they are restricted to non-negative integers in atomic games.

The expected cost that a flow of class c experiences for each (arbitrarily small) unitary fraction assigned to resource r becomes equal to

$$\alpha_r\left(n_r^{<c}(\boldsymbol{\sigma})+\frac{1}{n_r^c(\boldsymbol{\sigma})}\int_0^{n_r^c(\boldsymbol{\sigma})}t\,dt\right)+\beta_r=\alpha_r\left(n_r^{<c}(\boldsymbol{\sigma})+\frac{n_r^c(\boldsymbol{\sigma})}{2}\right)+\beta_r,$$

while the social cost in $\boldsymbol{\sigma}$ becomes

$$\mathsf{SC}(\boldsymbol{\sigma})=\sum_{r\in R}\left(\alpha_r\frac{n_r(\boldsymbol{\sigma})^2}{2}+\beta_r n_r(\boldsymbol{\sigma})\right).$$

The notion of pure Nash equilibrium assumes the following definition.

[4] We recall that notation $\boldsymbol{\sigma}_{-i},S$ denotes the strategy profile obtained from $\boldsymbol{\sigma}$ when player i changes her strategy from σ_i to S.

Definition 2. A strategy profile $\boldsymbol{\sigma}$ is a pure Nash equilibrium for Γ_p^{na} if and only if, for each $i \in [n]$, $j, k \in [m_i]$ with $\sigma_{i,j} > 0$, and $\delta \geq 0$,

$$\sum_{r \in S_{i,j} \setminus S_{i,k}} \left(\alpha_r \left(n_r^{<c(i)}(\boldsymbol{\sigma}) + \frac{n_r^{c(i)}(\boldsymbol{\sigma})}{2} \right) + \beta_r \right) \leq \sum_{r \in S_{i,k} \setminus S_{i,j}} \left(\alpha_r \left(n_r^{<c(i)}(\boldsymbol{\sigma}) + \frac{n_r^{c(i)}(\boldsymbol{\sigma}) + \delta}{2} \right) + \beta_r \right).$$

Letting δ go to zero, by continuity and monotonicity of the latency functions, the following characterization of pure Nash equilibria can be derived.

Lemma 1. *The strategy profile $\boldsymbol{\sigma}$ is a pure Nash equilibrium for Γ_p^{na} if and only if, for each $i \in [n]$ and $j, k \in [m_i]$ with $\sigma_{i,j} > 0$,*

$$\sum_{r \in S_{i,j}} \left(\alpha_r \left(n_r^{<c(i)}(\boldsymbol{\sigma}) + \frac{n_r^{c(i)}(\boldsymbol{\sigma})}{2} \right) + \beta_r \right) \leq \sum_{r \in S_{i,k}} \left(\alpha_r \left(n_r^{<c(i)}(\boldsymbol{\sigma}) + \frac{n_r^{c(i)}(\boldsymbol{\sigma})}{2} \right) + \beta_r \right).$$

Thus, if $\boldsymbol{\sigma}$ is a pure Nash equilibrium and $\boldsymbol{\sigma}^*$ is a social optimum for Γ_p^{na}, the following inequality, denoted as $ineq(i, j, k)$, holds for each $i \in [n]$, $j \in [m_i] : \sigma_{i,j} > 0$ and $k \in [m_i]$:

$$\sum_{r \in S_{i,j}} \left(\alpha_r \left(n_r^{<c(i)}(\boldsymbol{\sigma}) + \frac{n_r^{c(i)}(\boldsymbol{\sigma})}{2} \right) + \beta_r \right) - \sum_{r \in S_{i,k}} \left(\alpha_r \left(n_r^{<c(i)}(\boldsymbol{\sigma}) + \frac{n_r^{c(i)}(\boldsymbol{\sigma})}{2} \right) + \beta_r \right) \leq 0.$$

Multiplying $ineq(i, j, k)$ by $\frac{\sigma_{i,j}\sigma^*_{i,k}}{f_i}$ and then summing them up for any $j \in [m_i] : \sigma_{i,j} > 0$ and $k \in [m_i]$, we get

$$\sum_{j \in [m_i]: \sigma_{i,j} > 0} \frac{\sigma_{i,j}}{f_i} \sum_{r \in S_{i,j}} \left(\alpha_r \left(n_r^{<c(i)}(\boldsymbol{\sigma}) + \frac{n_r^{c(i)}(\boldsymbol{\sigma})}{2} \right) + \beta_r \right) \sum_{k \in [m_i]} \sigma^*_{i,k} - \sum_{k \in [m_i]} \frac{\sigma^*_{i,k}}{f_i} \sum_{r \in S_{i,k}} \left(\alpha_r \left(n_r^{<c(i)}(\boldsymbol{\sigma}) + \frac{n_r^{c(i)}(\boldsymbol{\sigma})}{2} \right) + \beta_r \right) \sum_{j \in [m_i]: \sigma_{i,j} > 0} \sigma_{i,j} \leq 0,$$

which, by using that $\sum_{j \in [m_i]: \sigma_{i,j} > 0} \sigma_{i,j} = \sum_{k \in [m_i]} \sigma^*_{i,k} = f_i$, yields

$$\sum_{j \in [m_i]: \sigma_{i,j} > 0} \sigma_{i,j} \sum_{r \in S_{i,j}} \left(\alpha_r \left(n_r^{<c(i)}(\boldsymbol{\sigma}) + \frac{n_r^{c(i)}(\boldsymbol{\sigma})}{2} \right) + \beta_r \right) - \sum_{k \in [m_i]} \sigma^*_{i,k} \sum_{r \in S_{i,k}} \left(\alpha_r \left(n_r^{<c(i)}(\boldsymbol{\sigma}) + \frac{n_r^{c(i)}(\boldsymbol{\sigma})}{2} \right) + \beta_r \right) \leq 0. \tag{2}$$

Moreover, for a pure Nash equilibrium $\boldsymbol{\sigma}$, by Definition 2, for every $i \in [n]$ and $j, j' \in [m_i]$ such that $\sigma_{i,j} > 0$ and $\sigma_{i,j'} > 0$, it must be

$$\sum_{r \in S_{i,j}} \left(\alpha_r \left(n_r^{<c(i)}(\boldsymbol{\sigma}) + \frac{n_r^{c(i)}(\boldsymbol{\sigma})}{2} \right) + \beta_r \right) = \sum_{r \in S_{i,j'}} \left(\alpha_r \left(n_r^{<c(i)}(\boldsymbol{\sigma}) + \frac{n_r^{c(i)}(\boldsymbol{\sigma})}{2} \right) + \beta_r \right) := C_i(\boldsymbol{\sigma}).$$

Thus, we get:

Lemma 2. *If $\boldsymbol{\sigma}$ is a pure Nash equilibrium for Γ_p^{na}, then $\mathsf{SC}(\boldsymbol{\sigma}) = \sum_{i\in[n]} C_i(\boldsymbol{\sigma}) f_i$.*

2.3. Price of anarchy and price of stability

Given a game Γ, either atomic or non-atomic, denote by $\mathsf{NE}(\Gamma)$ the set of its pure Nash equilibria. The price of anarchy of Γ is $\mathsf{PoA}(\Gamma) = \max_{\boldsymbol{\sigma}\in\mathsf{NE}(\Gamma)} \frac{\mathsf{SC}(\boldsymbol{\sigma})}{\mathsf{SC}(\boldsymbol{\sigma}^*)}$, while the price of stability of Γ is $\mathsf{PoS}(\Gamma) = \min_{\boldsymbol{\sigma}\in\mathsf{NE}(\Gamma)} \frac{\mathsf{SC}(\boldsymbol{\sigma})}{\mathsf{SC}(\boldsymbol{\sigma}^*)}$. For the priority-free case of $p = 1$, we have the following known results. For atomic games, [50] shows that the price of anarchy is $5/3$, while [13] proves that the price of stability drops to $1 + 1/\sqrt{5} \approx 1.447$. For non-atomic games, it is not difficult to see that both the random and the round-robin policy induce the same set of pure Nash equilibria. Hence, by the results in [54], both the price of anarchy and the price of stability are equal to $4/3$ (for instance, the classical Pigou's network yields a $4/3$ lower bound on the price of anarchy also in the random model).

3. Existence of pure Nash equilibria

In this section, we shall prove that priority-based congestion games always admit pure Nash equilibria. For non-atomic games, we also show that all equilibria attain the same social cost, thus implying that the price of anarchy and the price of stability coincide within this class.

3.1. Atomic games

A priority-based (affine) atomic congestion game with only one priority class boils down to a traditional congestion game for which existence of pure Nash equilibria (and more generally the finite improvement path property) is guaranteed by Rosenthal's Theorem [51]. However, for more priority classes, this equivalence does not hold any more and a dedicated existential proof is required.

Towards this end, we need to introduce some additional notation. Given an atomic game Γ_p^a, with $p \geq 1$, and a priority class $c \in [p]$, denote by $\Gamma_{\leq c}^a$ the restriction of Γ_p^a to the players of priority class at most c; moreover, given a strategy profile $\boldsymbol{\sigma}^{<c}$ for $\Gamma_{\leq c-1}^a$, with $\boldsymbol{\sigma}^0 := \emptyset$, denote by $\overline{\Gamma}_c^a(\boldsymbol{\sigma}^{<c})$ the game obtained from $\Gamma_{\leq c}^a$ by freezing the strategic choices of all players of class $c' < c$ according to $\boldsymbol{\sigma}^{<c}$ and letting only the players of class c play. We shall denote by $\overline{\boldsymbol{\sigma}}^c$ a strategy profile for $\overline{\Gamma}_c^a(\boldsymbol{\sigma}^{<c})$, i.e., a strategy profile satisfying $\overline{\sigma}_i^c = \sigma_i^{<c}$ for each player i such that $c(i) < c$.

We first show that, for any strategy profile $\boldsymbol{\sigma}^{<c}$ for $\Gamma_{\leq c-1}^a$, $\overline{\Gamma}_c^a(\boldsymbol{\sigma}^{<c})$ is an exact potential game.

Lemma 3. *For any $p \geq 2$, affine atomic congestion game with p priority classes Γ_p^a, priority class $c \in [p]$ and strategy profile $\boldsymbol{\sigma}^{<c}$ for $\Gamma_{\leq c-1}^a$, game $\overline{\Gamma}_c^a(\boldsymbol{\sigma}^{<c})$ admits the following exact potential function[5]:*

$$\Phi_c(\overline{\boldsymbol{\sigma}}^c) := \sum_{r\in R} \left(\alpha_r n_r^c(\overline{\boldsymbol{\sigma}}^c) \left(n_r^{<c}(\overline{\boldsymbol{\sigma}}^c) + \frac{n_r^c(\boldsymbol{\sigma}) + 3}{4} \right) + n_r^c(\overline{\boldsymbol{\sigma}}^c)\beta_r \right). \tag{3}$$

Proof. Fix an integer $p \geq 2$, an affine atomic congestion game with p priority classes Γ_p^a, a priority class $c \in [p]$, a strategy profile $\boldsymbol{\sigma}^{<c}$ for $\Gamma_{\leq c-1}^a$, a strategy profile $\overline{\boldsymbol{\sigma}}^c$ for $\overline{\Gamma}_c^a(\boldsymbol{\sigma}^{<c})$, a player i and a strategy $t \in \mathcal{S}_i$. Observe that, by definition, it must be $c(i) = c$. We shall prove $cost_i(\overline{\boldsymbol{\sigma}}^c) - cost_i(\overline{\boldsymbol{\sigma}}_{-i}^c, t) = \Phi_c(\overline{\boldsymbol{\sigma}}^c) - \Phi_c(\overline{\boldsymbol{\sigma}}_{-i}^c, t)$. We have

$$\begin{aligned}
&\Phi_c(\overline{\boldsymbol{\sigma}}^c) - \Phi_c(\overline{\boldsymbol{\sigma}}_{-i}^c, t)\\
&= \sum_{r\in\sigma_i\setminus t} \left(\alpha_r n_r^c(\overline{\boldsymbol{\sigma}}^c) \left(n_r^{<c}(\overline{\boldsymbol{\sigma}}^c) + \frac{n_r^c(\overline{\boldsymbol{\sigma}}^c) + 3}{4} \right) + n_r^c(\overline{\boldsymbol{\sigma}}^c)\beta_r \right)\\
&\quad - \sum_{r\in\sigma_i\setminus t} \left(\alpha_r \left(n_r^c(\overline{\boldsymbol{\sigma}}^c) - 1 \right) \left(n_r^{<c}(\overline{\boldsymbol{\sigma}}^c) + \frac{n_r^c(\overline{\boldsymbol{\sigma}}^c) + 2}{4} \right) + \left(n_r^c(\overline{\boldsymbol{\sigma}}^c) - 1 \right)\beta_r \right)\\
&\quad + \sum_{r\in t\setminus\sigma_i} \left(\alpha_r n_r^c(\overline{\boldsymbol{\sigma}}^c) \left(n_r^{<c}(\overline{\boldsymbol{\sigma}}^c) + \frac{n_r^c(\overline{\boldsymbol{\sigma}}^c) + 3}{4} \right) + n_r^c(\overline{\boldsymbol{\sigma}}^c)\beta_r \right)\\
&\quad - \sum_{r\in t\setminus\sigma_i} \left(\alpha_r \left(n_r^c(\overline{\boldsymbol{\sigma}}^c) + 1 \right) \left(n_r^{<c}(\overline{\boldsymbol{\sigma}}^c) + \frac{n_r^c(\overline{\boldsymbol{\sigma}}^c) + 4}{4} \right) + \left(n_r^c(\overline{\boldsymbol{\sigma}}^c) + 1 \right)\beta_r \right)
\end{aligned}$$

[5] We recall that an exact potential function is a function mapping strategy profiles to the reals in such a way that the difference in the potential of two profiles differing for the choice of a unique player i equals the difference of the cost that i experiences in the two profiles. By definition, a game possessing an exact potential function admits pure Nash equilibria.

$$
\begin{aligned}
&= \sum_{r\in\sigma_i\setminus t}\left(\alpha_r\left(n_r^{<c}(\overline{\boldsymbol{\sigma}}^c)+\frac{n_r^c(\overline{\boldsymbol{\sigma}}^c)+1}{2}\right)+\beta_r\right)\\
&\quad - \sum_{r\in t\setminus\sigma_i}\left(\alpha_r\left(n_r^{<c}(\overline{\boldsymbol{\sigma}}^c)+\frac{n_r^c(\overline{\boldsymbol{\sigma}}^c)+2}{2}\right)+\beta_r\right)\\
&= cost_i(\overline{\boldsymbol{\sigma}}^c) - cost_i(\overline{\boldsymbol{\sigma}}^c_{-i}, t). \quad \square
\end{aligned}
$$

We can now prove that any priority-based affine atomic congestion game admits a pure Nash equilibrium.

Theorem 1. *For any $p \geq 1$, game Γ_p^a admits pure Nash equilibria.*

Proof. As the cost of a player of priority class c is not influenced by the choices of the players of higher classes, it follows that, given a strategy profile $\boldsymbol{\sigma}$ for Γ_p^a and a player i, it holds that σ_i is a best-response for i against $\boldsymbol{\sigma}_{-i}$ in Γ_p^a if and only if σ_i is a best-response for i against $\overline{\boldsymbol{\sigma}}_{-i}^{c(i)}$ in $\overline{\Gamma}_{c(i)}^a(\boldsymbol{\sigma}^{<c(i)})$. Thus, for each $c \in [p]$, thanks to Lemma 3, a pure Nash equilibrium for $\Gamma_{\leq c}^a$ can be constructed inductively by extending a pure Nash equilibrium for $\Gamma_{\leq c-1}^a$. $\square$

3.2. Non-atomic games

By following and extending [54], we show that every non-atomic priority-based affine congestion game admits pure Nash equilibria and that there is no difference between the price of anarchy and the price of stability within this class.

Theorem 2. *Every non-atomic priority-based affine congestion game admits pure Nash equilibria. Moreover, all equilibria have the same social cost.*

Proof. Given a non-atomic game Γ_p^{na}, with $p \geq 1$, and a priority class $c \in [p]$, denote by $\Gamma_{\leq c}^{na}$ the restriction of Γ_p^{na} to the flows of priority class at most c; moreover, given a strategy profile $\boldsymbol{\sigma}^{<c}$ for $\Gamma_{\leq c-1}^{na}$, with $\boldsymbol{\sigma}^0 := \emptyset$, denote by $\overline{\Gamma}_c^{na}(\boldsymbol{\sigma}^{<c})$ the game obtained from $\Gamma_{\leq c}^{na}$ by freezing the strategic choices of all flows of class $c' < c$ according to $\boldsymbol{\sigma}^{<c}$ and letting only the flows of class c play. We shall denote by $\overline{\boldsymbol{\sigma}}^c$ a strategy profile for $\overline{\Gamma}_c^{na}(\boldsymbol{\sigma}^{<c})$, i.e., a strategy profile satisfying $\overline{\sigma}_{i,j}^c = \sigma_{i,j}^{<c}$ for each flow i such that $c(i) < c$ and $j \in [m_i]$.

We first show that, for any strategy profile $\boldsymbol{\sigma}^{<c}$ for $\Gamma_{\leq c-1}^{na}$, $\overline{\Gamma}_c^{na}(\boldsymbol{\sigma}^{<c})$ always admits pure Nash equilibria, all having the same social cost.

Fix a non-atomic priority-based affine congestion game Γ_p^{na} and a priority class $c \in [p]$. For every resource $r \in R$, define $h_r(x) = \int_0^x \left(\alpha_r\left(n_r(\boldsymbol{\sigma}^{<c}) + \frac{t}{2}\right) + \beta_r\right) dt$. As h_r is differentiable and has a non-decreasing derivative, h_r is convex for each $r \in R$. Consider the following mathematical program, denoted as MP:

$$
\begin{aligned}
\min \ & \sum_{r\in R} h_r(x)\\
s.t. \ & \sum_{j\in[m_i]} \sigma_{i,j} = f_i, \quad \forall i\in[n] : c(i) = c\\
& x = \sum_{i\in[n]:c(i)=c}\ \sum_{j\in[m_i]:r\in S_{i,j}} \sigma_{i,j}, \quad \forall r \in R,\\
& \sigma_{i,j} \geq 0, \quad \forall i \in [n] : c(i) = c,\ \forall j \in [m_i]
\end{aligned}
$$

As the objective function of MP is convex and so is also its feasible set, by applying the first order condition for optimality in convex optimization, we get that a strategy profile $\overline{\boldsymbol{\sigma}}^c$ is an optimal solution for MP if and only if $\sum_{r\in R} \nabla h_r(x)^T(\boldsymbol{\sigma}' - \overline{\boldsymbol{\sigma}}^c) \geq 0$ for each strategy profile $\boldsymbol{\sigma}' \in \overline{\Gamma}_c^{na}(\boldsymbol{\sigma}^{<c})$. As $\nabla h_r(x) = \alpha_r\left(n_r(\boldsymbol{\sigma}^{<c}) + \frac{x}{2}\right) + \beta_r$, it immediately follows that $\overline{\boldsymbol{\sigma}}^c$ is an optimal solution for MP if and only if for each $i \in [n]$ and $j,k \in [m_i]$ such that $\sigma_{i,j} > 0$ (by choosing $\boldsymbol{\sigma}'$ such that $\boldsymbol{\sigma}'$ is obtained from $\overline{\boldsymbol{\sigma}}^c$ by moving an infinitesimal amount of flow from $S_{i,j}$ to $S_{i,k}$), it holds that

$$
\begin{aligned}
&\sum_{r\in S_{i,j}}\left(\alpha_r\left(n_r^{<c}(\overline{\boldsymbol{\sigma}}^c)+\frac{n_r^c(\overline{\boldsymbol{\sigma}}^c)}{2}\right)+\beta_r\right) \leq\\
&\quad \sum_{r\in S_{i,k}}\left(\alpha_r\left(n_r^{<c}(\overline{\boldsymbol{\sigma}}^c)+\frac{n_r^c(\overline{\boldsymbol{\sigma}}^c)}{2}\right)+\beta_r\right),
\end{aligned}
$$

that is, if and only if $\overline{\boldsymbol{\sigma}}^c$ is a pure Nash equilibrium for $\overline{\Gamma}_c^{na}(\boldsymbol{\sigma}^{<c})$ (see Lemma 1). As the feasible set of MP is trivially non-empty, this proves existence of pure Nash equilibria. Now, assume that $\overline{\boldsymbol{\sigma}}^c$ and $\boldsymbol{\sigma}'$ are pure Nash equilibria for $\overline{\Gamma}_c^{na}(\boldsymbol{\sigma}^{<c})$. By the convexity of the objective function of MP, whenever $n_r^c(\overline{\boldsymbol{\sigma}}^c) \neq n_r^c(\boldsymbol{\sigma}')$, function h_r must be linear between these two values, otherwise a feasible solution for MP with a lower objective value can be derived through a convex combination of $\overline{\boldsymbol{\sigma}}^c$ and $\boldsymbol{\sigma}'$. This implies

$$\sum_{r\in S_{i,j}}\left(\alpha_r\left(n_r^{<c}(\overline{\boldsymbol{\sigma}}^c)+\frac{n_r^c(\overline{\boldsymbol{\sigma}}^c)}{2}\right)+\beta_r\right)= \sum_{r\in S_{i,j}}\left(\alpha_r\left(n_r^{<c}(\boldsymbol{\sigma}')+\frac{n_r^c(\boldsymbol{\sigma}')}{2}\right)+\beta_r\right)$$

for each $i\in[n]:c(i)=c$ and $j\in[m_i]$, which yields $C_i(\overline{\boldsymbol{\sigma}}^c)=C_i(\boldsymbol{\sigma}')$ for each $i\in[n]:c(i)=c$. By Lemma 2, $\mathsf{SC}(\boldsymbol{\sigma})=\mathsf{SC}(\overline{\boldsymbol{\sigma}})$ follows.

By using the same inductive argument exploited in the proof of Theorem 1, the claimed result for game Γ_p^{na} follows. □

4. Bounding the price of anarchy and the price of stability

In this section, we characterize the price of anarchy and the price of stability of priority-based affine congestion games for both of their versions: atomic and non-atomic. We perform our analysis by relying on the primal-dual method introduced in [6]. As the base case of $p=1$ has already been solved, we shall focus on games with at least two priority classes.

4.1. The primal-dual formulation

Fix a priority-based affine congestion game Γ_p with $p\geq 2$, a pure Nash equilibrium $\boldsymbol{\sigma}$ for Γ_p and a social optimum $\boldsymbol{\sigma}^*$ for Γ_p. For a resource $r\in R$ and a priority class $c\in[p]$, set $k_r^c=n_r^c(\boldsymbol{\sigma})$, $o_r^c=n_r^c(\boldsymbol{\sigma}^*)$, $k_r^{<c}=n_r^{<c}(\boldsymbol{\sigma})$, $o_r^{<c}=n_r^{<c}(\boldsymbol{\sigma}^*)$, $k_r=n_r(\boldsymbol{\sigma})$ and $o_r=n_r(\boldsymbol{\sigma}^*)$. Observe that, no matter whether Γ_p is an atomic or non-atomic game, all the previous quantities are well defined. We recall that these values are non-negative reals in non-atomic games and non-negative integers in atomic ones.

If Γ_p is an atomic game, for each $c\in[p]$, by summing inequality (1), derived from the definition of pure Nash equilibria for atomic games (Definition 1), for each $i\in[n]$ such that $c(i)=c$, we obtain

$$\sum_{r\in R}\left(\alpha_r\left(k_r^c\left(k_r^{<c}+\frac{k_r^c+1}{2}\right)-o_r^c\left(k_r^{<c}+\frac{k_r^c+2}{2}\right)\right)+\beta_r\left(k_r^c-o_r^c\right)\right)\leq 0.$$

Similarly, if Γ_p is a non-atomic game, for each $c\in[p]$, by summing inequality (2), derived from the definition of pure Nash equilibria for non-atomic games (Definition 2), for each $i\in[n]$ such that $c(i)=c$, we get

$$\sum_{r\in R}\left(\alpha_r\left(k_r^c\left(k_r^{<c}+\frac{k_r^c}{2}\right)-o_r^c\left(k_r^{<c}+\frac{k_r^c}{2}\right)\right)+\beta_r\left(k_r^c-o_r^c\right)\right)\leq 0.$$

Thus, using an auxiliary variable δ such that $\delta=1$ when dealing with atomic games and $\delta=0$ when dealing with non-atomic ones, we have that inequality

$$\sum_{r\in R}\left(\alpha_r\left(k_r^c\left(k_r^{<c}+\frac{k_r^c+\delta}{2}\right)-o_r^c\left(k_r^{<c}+\frac{k_r^c+2\delta}{2}\right)\right)\right) +\sum_{r\in R}\left(\beta_r\left(k_r^c-o_r^c\right)\right)\leq 0 \tag{4}$$

holds for each $c\in[p]$. Moreover, also the social cost of both $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma}^*$ can be expressed in a unified manner, as we have

$$\mathsf{SC}(\boldsymbol{\sigma})=\sum_{r\in R}\left(\alpha_r\frac{k_r(k_r+\delta)}{2}+\beta_r k_r\right)\text{ and}$$
$$\mathsf{SC}(\boldsymbol{\sigma}^*)=\sum_{r\in R}\left(\alpha_r\frac{o_r(o_r+\delta)}{2}+\beta_r o_r\right)$$

with the same constraints on δ.

The primal-dual method bounds the price of anarchy of a game by choosing latency functions maximizing $\mathsf{SC}(\boldsymbol{\sigma})$ under the constraints that $\mathsf{SC}(\boldsymbol{\sigma}^*)=1$ and the pair $(\boldsymbol{\sigma},\boldsymbol{\sigma}^*)$ satisfies inequality (4). By applying the method to Γ_p, we get the following primal linear program $PP(\Gamma_p)$:

$$\begin{aligned}
\max\ & \sum_{r\in R}\left(\alpha_r \frac{k_r(k_r+\delta)}{2}+\beta_r k_r\right)\\
s.t.\ & \sum_{r\in R}\left(\alpha_r\left(k_r^c\left(k_r^{<c}+\frac{k_r^c+\delta}{2}\right)-o_r^c\left(k_r^{<c}+\frac{k_r^c+2\delta}{2}\right)\right)\right)\\
&\qquad +\sum_{r\in R}\left(\beta_r\left(k_r^c-o_r^c\right)\right)\le 0,\quad \forall c\in[p]\\
& \sum_{r\in R}\left(\alpha_r \frac{o_r(o_r+\delta)}{2}+\beta_r o_r\right)=1\\
& \alpha_r,\beta_r\ge 0,\quad \forall r\in R.
\end{aligned}$$

The dual program $DP(\Gamma_p)$, obtained by associating dual variable x_c to each of the first p constraints of $PP(\Gamma_p)$ and dual variable γ to the last constraint of $PP(\Gamma_p)$, is the following:

$$\begin{aligned}
\min\ & \gamma\\
s.t.\ & \sum_{c\in[p]}\left(x_c\left(k_r^c\left(k_r^{<c}+\frac{k_r^c+\delta}{2}\right)-o_r^c\left(k_r^{<c}+\frac{k_r^c+2\delta}{2}\right)\right)\right)\\
&\ +\gamma\frac{o_r(o_r+\delta)}{2}-\frac{k_r(k_r+\delta)}{2}\ge 0\qquad \forall r\in R &(5)\\
& \sum_{c\in[p]}\left(x_c(k_r^c-o_r^c)\right)+\gamma o_r-k_r\ge 0\qquad \forall r\in R &(6)\\
& x_c\ge 0\qquad \forall c\in[p].
\end{aligned}$$

4.2. Solving the dual program

Our aim is to determine an optimal solution for $DP(\Gamma_p)$, for any $p\ge 2$. To this end, we make use of an auxiliary non-linear program, denoted as $NLP(p)$, and defined as follows:

$$\begin{aligned}
\min\ & \gamma\\
s.t.\ & x_1\le\gamma &(7)\\
& x_{c+1}^2\le\gamma(x_c-1)\quad \forall c\in[p-2] &(8)\\
& x_{p-1}\le\gamma(x_{p-1}-1) &(9)\\
& x_p=\frac{2x_{p-1}}{x_{p-1}+1} &(10)\\
& x_c\ge 0\qquad \forall c\in[p]. &(11)
\end{aligned}$$

Observation 1. For the sake of readability, we briefly explain how $NLP(p)$ has been derived. Constraint (7) is obtained from constraint (5) when $\delta=1$, by choosing $o_r^1=o_r=1$ and setting to zero all other terms. The same constraint can be obtained also in the case of $\delta=0$ by means of the same settings, but applied to constraint (6). For each $c\in[p-2]$, constraint (8) is obtained, independently of the value of δ, from constraint (5) by choosing $k_r^c=k_r=\theta o_r$, $o_r^{c+1}=o_r$, setting to zero all other terms, letting o_r go to infinity, and then optimizing over θ. Constraint (9) is obtained, independently of the value of δ, from constraint (5) by choosing $k_r^{p-1}=\theta o_r$, $k_r^p=o_r^p$, $k_r=k_r^{p-1}+k_r^p$, $o_r^p=o_r$, setting to zero all other terms, letting o_r go to infinity, and then optimizing over θ. The value of x_p does not influence the value of the objective function, so constraint (10) could be indeed arbitrary, as long as the value of x_p satisfies the technical conditions needed in our later proofs.

As our main result, we show that $NLP(p)$ admits a unique optimal solution which is also feasible for $DP(\Gamma_p)$.

Theorem 3. *For every $p\ge 2$, there exists a unique optimal solution $\overline{\mathbf{s}}(p)=(\overline{x}_1(p),\dots,\overline{x}_p(p),\overline{\gamma}(p))$ for $NLP(p)$. Moreover, $\overline{\mathbf{s}}(p)$ is a feasible solution for $DP(\Gamma_p)$.*

We proceed to proving the first part of Theorem 3 through a sequence of steps, while the second claim will be shown in the next subsection. The following lemma provides a characterization of an optimal solution for $NLP(p)$.

Lemma 4. *For every $p \geq 2$, there exists a unique optimal solution $\overline{\boldsymbol{s}}(p) = (\overline{x}_1(p), \ldots, \overline{x}_p(p), \overline{\gamma}(p))$ for $NLP(p)$ satisfying constraints (7)–(9) at equality and such that (i) $\overline{x}_c(p) > \overline{x}_{c+1}(p)$ for each $c \in [p-1]$, (ii) $\overline{x}_c(p) \geq 4/3$ for each $c \in [p-1]$, (iii) $8/7 \leq \overline{x}_p(p) \leq 8/5$, and (iv) $\overline{x}_1(p) = 2$ for $p = 2$ and $\overline{x}_1(p) \geq 1 + 2/\sqrt{3}$ for every $p \geq 3$.*

Proof. Fix a value $p \geq 2$ and observe that $NLP(p)$ admits a non-empty set of feasible solutions (for instance, it suffices setting $x_c = 2$ for every $c \in [p-1]$, $x_p = 4/3$ and $\gamma = 4$).

First of all, we show that, in any feasible solution for $NLP(p)$, it must be $x_c > 1$ for each $c \in [p]$. Note that by constraints (7) and (11), we get $\gamma \geq 0$. For each $c \in [p-2]$, by constraints (8) and (11), it follows $x_c \geq 1$, otherwise $\gamma \geq 0$ is contradicted. Similarly, by constraint (9), we get $x_{p-1} \geq 1$. So, we conclude that $x_c \geq 1$ for each $c \in [p-1]$. Assume that $x_c = 1$ for some $c \in [p-2]$. By constraint (8), we get $x_{c+1} = 0$, thus contradicting $x_c \geq 1$ for each $c \in [p-1]$. Also, $x_{p-1} = 1$ contradicts constraint (9). So, $x_c > 1$ for each $c \in [p-1]$. Given that $x_{p-1}(p) > 1$, by constraint (10), it follows $x_p(p) > 1$.

Now note that, for each $c \in [p-1]$, $NLP(p)$ has exactly two constraints, among the set of constraints (7)–(9), involving variable x_c which give an upper and a lower bound on γ, respectively. This is sufficient to conclude that there exists a unique optimal solution $\overline{\boldsymbol{s}}(p) = (\overline{x}_1(p), \ldots, \overline{x}_p(p), \overline{\gamma}(p))$ satisfying constraints (7)–(9) at equality.

Using this property, we show claim *(i)* by induction. For the base case of $p = 2$, we have that constraint (8), which rewrites as $\overline{x}_1(p)(\overline{x}_1(p) - 1) = \overline{x}_2(p)^2$, implies $\overline{x}_1(p) > \overline{x}_2(p)$. Now assume the claim true up to $c - 1 \leq p - 3$. As $c \leq p - 2$, by constraint (8), we have $\frac{\overline{x}_c(p)^2}{\overline{x}_{c-1}(p)-1} = \frac{\overline{x}_{c+1}(p)^2}{\overline{x}_c(p)-1}$ which implies $\overline{x}_c(p)^2(\overline{x}_c(p) - 1) = \overline{x}_{c+1}(p)^2(\overline{x}_{c-1}(p) - 1)$. As, by the inductive hypothesis, we have $\overline{x}_{c-1}(p) > \overline{x}_c(p)$, by using $\overline{x}_c(p) > 1$, it follows that $\overline{x}_c(p)^2 > \overline{x}_{c+1}(p)^2$, which shows $\overline{x}_c(p) > \overline{x}_{c+1}(p)$ for each $c \in [p-2]$. By constraint (10), using $\overline{x}_{p-1}(p) > 1$ we get $\overline{x}_{p-1}(p) > \overline{x}_p(p)$, thus showing claim *(i)*.

As $\overline{\boldsymbol{s}}(p)$ is an optimal solution and we have shown the existence of a feasible solution with $\gamma = 4$, it must be $\overline{\gamma}(p) \leq 4$ and, by constraint (7), $\overline{x}_1(p) \leq 4$. By constraint (9), it must be $\overline{x}_{p-1}(p) \geq 4/3$ which, by claim *(i)*, implies $\overline{x}_c(p) \geq 4/3$ for each $c \in [p-1]$, that is, claim *(ii)*. Moreover, as $\overline{x}_{p-1}(p) \geq 4/3$ and $\overline{x}_{p-1}(p) \leq \overline{x}_1(p) \leq \overline{\gamma}(p) \leq 4$, by constraint (10), claim (*iii*) follows.

To prove claim *(iv)*, observe that there exists a unique optimal solution $\overline{\boldsymbol{s}}(2) = (\overline{x}_1(2), \overline{x}_2(2), \overline{\gamma}(2)) = (2, 4/3, 2)$ for $NLP(2)$ yielding $\overline{\gamma}(2) = \overline{x}_1(2) = 2$. Moreover, there exists a unique optimal solution $\overline{\boldsymbol{s}}(3) = (\overline{x}_1(3), \overline{x}_2(3), \overline{x}_3(3), \overline{\gamma}(3))$ for $NLP(3)$ such that $\overline{x}_1(3) = \overline{\gamma}(3) \approx 2.3247 > 1 + 2/\sqrt{3}$. As p increases, the set of feasible solutions for $NLP(p)$ shrinks, which implies that $\overline{\gamma}(p)$ increases with p, i.e., $\overline{\gamma}(p) \geq 2.3248$. Since we have shown that $\overline{\boldsymbol{s}}(p)$ makes constraint (7) tight, claim *(iv)* follows. $\square$

More properties of the optimal solution $\overline{\boldsymbol{s}}(p)$ are given in the following.

Lemma 5. *For each $p \geq 2$, we have $1 + \overline{x}_{p-1}(p) - 2\overline{x}_p(p) > 0$ and $1 + \overline{x}_{c-1}(p) - 2\overline{x}_c(p) \leq 0$ for each $2 \leq c < p$.*

Proof. Fix a value $p \geq 2$. By constraint (10) of $NLP(p)$, we get $1 + \overline{x}_{p-1}(p) - 2\overline{x}_p(p) = 1 + \overline{x}_{p-1}(p) - \frac{4\overline{x}_{p-1}(p)}{\overline{x}_{p-1}(p)+1} = \frac{(\overline{x}_{p-1}(p)-1)^2}{\overline{x}_{p-1}(p)+1} > 0$ as $\overline{x}_{p-1}(p) \geq 4/3$ by Lemma 4. For any $c \in [p-1] \setminus \{1\}$, by constraint (8) of NLP_p, we get $1 + \overline{x}_{c-1}(p) - 2\overline{x}_c(p) = 1 + \overline{x}_{c-1}(p) - 2\sqrt{\overline{\gamma}(p)(\overline{x}_{c-1}(p) - 1)} \leq 0$ if and only if it holds that

$$\overline{\gamma}(p) \geq \frac{(\overline{x}_{c-1}(p) + 1)^2}{4(\overline{x}_{c-1}(p) - 1)}. \tag{12}$$

It is easy to check that, as $\overline{x}_{c-1}(p) > 1$ by Lemma 4, the right-hand side of inequality (12) is maximized when $\overline{x}_{c-1}(p) \in \{\overline{x}_1(p), \overline{x}_{p-2}(p)\}$.

Assume first that $\overline{x}_{c-1}(p) = \overline{x}_1(p)$ which requires $p \geq 3$. By constraint (7) of $NLP(p)$, inequality (12) boils down to $\overline{x}_1(p) \geq \frac{(\overline{x}_1(p)+1)^2}{4(\overline{x}_1(p)-1)}$ which is always satisfied as long as $\overline{x}_1(p) \geq 1 + 2/\sqrt{3} \approx 2.1547$ which holds by claim (*iv*) of Lemma 4. Now assume $\overline{x}_{c-1}(p) = \overline{x}_{p-2}(p)$ which again requires $p \geq 3$. By constraints (9) and (8) of $NLP(p)$, we have $\overline{x}_{p-1}(p) = \overline{\gamma}(p)/(\overline{\gamma}(p) - 1)$ and $\overline{\gamma}(p) = \frac{\overline{x}_{p-1}(p)^2}{\overline{x}_{p-2}(p)-1}$ which combined together give $\overline{x}_{p-2}(p) = \frac{\overline{\gamma}(p)}{(\overline{\gamma}(p)-1)^2} + 1$. By using this equality in inequality (12), we obtain $\overline{\gamma}(p) \geq \frac{(2\overline{\gamma}(p)^2 - 3\overline{\gamma}(p)+2)^2}{4\overline{\gamma}(p)(\overline{\gamma}(p)-1)^2}$ which is always satisfied whenever $\overline{\gamma}(p) \geq 2$. $\square$

4.3. Proof of feasibility

We now show that $\overline{\boldsymbol{s}}(p)$ is feasible for $DP(\Gamma_p)$. We start with the set of dual constraints defined by (6) which is easier to analyze. By claim *(i)* of Lemma 4 and by constraint (7) of $NLP(p)$, we get $\overline{\gamma}(p) \geq \overline{x}_1(p) > \ldots > \overline{x}_p(p) > 1$. This implies $\sum_{c\in[p]} \left(\overline{x}_c(p) k_r^c\right) \geq \overline{x}_p(p) \sum_{c\in[p]} k_r^c = \overline{x}_p(p) k_r \geq k_r$ for each $k_r \geq 0$ and $\sum_{c\in[p]} \left(\overline{x}_c(p) o_r^c\right) \leq \overline{x}_1(p) \sum_{c\in[p]} o_r^c = \overline{x}_1(p) o_r \leq \overline{\gamma}(p) o_r$ for each $o_r \geq 0$. Thus, any possible dual constraint defined by (6) is satisfied by $\overline{\boldsymbol{s}}(p)$.

Additional work is needed to handle the dual constraints defined by (5). Observe that, once the values of $\overline{\boldsymbol{s}}(p)$ are fixed, as the coefficients $k_r^{<c}$, k_r, $o_r^{<c}$ and o_r are all obtained as a function of the tuple $\left(k_r^c, o_r^c\right)_{c\in[p]}$, each constraint in (5) can be completely specified by a vector $\boldsymbol{v} \in \mathbb{V}$, with $\mathbb{V} := \mathbb{Z}_{\geq 0}^{2p}$ for atomic games and $\mathbb{V} := \mathbb{R}_{\geq 0}^{2p}$ for non-atomic ones. Thus, given a

vector $\boldsymbol{v} \in \mathbb{V}$, define $\psi(\boldsymbol{v})$ as the left-hand side of the dual constraint induced by $\boldsymbol{v}$. Let $\mathcal{C} = \{\psi(\boldsymbol{v}) \geq 0 : \boldsymbol{v} \in \mathbb{V}\}$ denote the set of all possible dual constraints defined by (5). Our aim is to determine a *representative subset of constraints* $\mathcal{C}' \subseteq \mathcal{C}$ such that showing that $\overline{\mathbf{s}}(p)$ satisfies $\mathcal{C}'$ will be sufficient to prove that $\overline{\mathbf{s}}(p)$ satisfies $\mathcal{C}$.

Within this subsection, we shall always consider a fixed resource $r \in R$ and a fixed integer $p \geq 2$. Thus, for the ease of simplicity, we shall drop the subscript r from the notation and the argument p from the elements of the optimal solution $\overline{\mathbf{s}}(p)$. Thus, we shall write k^c, $k^{<c}$, k, o^c, o and $\overline{x}_c$ in place of k_r^c, $k_r^{<c}$, k_r, o_r^c, o_r and $\overline{x}_c(p)$. For a priority class $c \in [p]$, define $\Delta^c = k^c - o^c$. We start by characterizing the structure of the values $(o^c)_{c\in[p]}$ in the dual constraints belonging to $\mathcal{C}'$. Say that a dual constraint is *homogeneous with respect to* $\boldsymbol{\sigma}^*$ if there exists a class c such that $o^c = o$.

Lemma 6. *If $\overline{\mathbf{s}}(p)$ satisfies all constraints in $\mathcal{C}$ that are homogeneous with respect to $\boldsymbol{\sigma}^*$, then $\overline{\mathbf{s}}(p)$ satisfies $\mathcal{C}$.*

Proof. Assume, by way of contradiction, that there is a constraint $\psi(\boldsymbol{v}) \geq 0 \in \mathcal{C}$ such that $o^c, o^{c'} > 0$ for some pair of distinct classes c and c', with $c < c'$, which is not satisfied by $\overline{\mathbf{s}}(p)$, i.e., we assume that $\psi(\boldsymbol{v}) < 0$ holds. Let ϵ be an arbitrary quantity such that $0 < \epsilon \leq \min\{o^c, o^{c'}\}$ with $\epsilon \in \mathbb{Z}$ in case of atomic games. The contribution of classes c and c' to $\psi(\boldsymbol{v})$ can be rewritten as

$$\overline{x}_c[\Delta^c(k^{<c} + (k^c + \delta)/2) - o^c\delta/2] + \overline{x}_{c'}[\Delta^{c'}(k^{<c'} + (k^{c'} + \delta)/2) - o^{c'}\delta/2].$$

Let $\overrightarrow{\boldsymbol{v}}$ be the tuple obtained from $\boldsymbol{v}$ by moving a quantity of ϵ from o^c to $o^{c'}$. The contribution of classes c and c' to $\psi(\overrightarrow{\boldsymbol{v}})$ is equal to

$$\begin{aligned}&\overline{x}_c[(\Delta^c + \epsilon)(k^{<c} + (k^c + \delta)/2) - (o^c - \epsilon)\delta/2] \\ &\qquad + \overline{x}_{c'}[(\Delta^{c'} - \epsilon)(k^{<c'} + (k^{c'} + \delta)/2) - (o^{c'} + \epsilon)\delta/2].\end{aligned}$$

Thus, we get $\psi(\boldsymbol{v}) - \psi(\overrightarrow{\boldsymbol{v}}) = -\overline{x}_c\epsilon[k^{<c} + k^c/2 + \delta] + \overline{x}_{c'}\epsilon[k^{<c'} + k^{c'}/2 + \delta]$.

Let $\overleftarrow{\boldsymbol{v}}$ be the tuple obtained from $\boldsymbol{v}$ by moving a quantity of ϵ from $o^{c'}$ to o^c. The contribution of classes c and c' to $\psi(\overleftarrow{\boldsymbol{v}})$ is equal to

$$\begin{aligned}&\overline{x}_c[(\Delta^c - \epsilon)(k^{<c} + (k^c + \delta)/2) - (o^c + \epsilon)\delta/2] \\ &\qquad + \overline{x}_{c'}[(\Delta^{c'} + \epsilon)(k^{<c'} + (k^{c'} + \delta)/2) - (o^{c'} - \epsilon)\delta/2].\end{aligned}$$

Thus, we get $\psi(\boldsymbol{v}) - \psi(\overleftarrow{\boldsymbol{v}}) = \overline{x}_c\epsilon[k^{<c} + k^c/2 + \delta] - \overline{x}_{c'}\epsilon[k^{<c'} + k^{c'}/2 + \delta]$.

As $\psi(\boldsymbol{v}) - \psi(\overrightarrow{\boldsymbol{v}}) = -\left(\psi(\boldsymbol{v}) - \psi(\overleftarrow{\boldsymbol{v}})\right)$, it follows that $\min\{\psi(\overrightarrow{\boldsymbol{v}}), \psi(\overleftarrow{\boldsymbol{v}})\} \leq \psi(\boldsymbol{v}) < 0$, which implies that one of the two transformations produces a dual constraint which is also not satisfied by $\overline{\mathbf{s}}(p)$. Since $|\psi(\boldsymbol{v}) - \psi(\overrightarrow{\boldsymbol{v}})|$ does not depend on both o^c and $o^{c'}$, the chosen transformation can be repeated until either $o^c = 0$ or $o^{c'} = 0$. By using this argument for all pairs c and c' such that $o^c, o^{c'} > 0$, one finally gets a constraint which is homogeneous with respect to $\boldsymbol{\sigma}^*$ and is not satisfied by $\overline{\mathbf{s}}(p)$: a contradiction. □

We now proceed by characterizing the structure of the values $(k^c)_{c\in[p]}$ in the constraints belonging to $\mathcal{C}'$. We say that a constraint is *almost homogeneous with respect to both* $\boldsymbol{\sigma}$ *and* $\boldsymbol{\sigma}^*$ if either *(i)* there exists a priority class c such that $o^c = o$ and $k^{c'} > 0$ if and only if $c' \in \{c-1, c\}$, or *(ii)* $o = 0$ and $k^{c'} > 0$ if and only if $c' = p$.

Lemma 7. *If $\overline{\mathbf{s}}(p)$ satisfies all constraints in $\mathcal{C}$ that are almost homogeneous with respect to both $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma}^*$, then $\overline{\mathbf{s}}(p)$ satisfies $\mathcal{C}$.*

Proof. Assume, by way of contradiction and by Lemma 6, that there exists a constraint $\psi(\boldsymbol{v}) \geq 0 \in \mathcal{C}$ which is homogeneous with respect to $\boldsymbol{\sigma}^*$, not almost homogeneous with respect to both $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma}^*$, and not satisfied by $\overline{\mathbf{s}}(p)$.

Let us consider first the case of $o > 0$, that is, $o^c = o$ for some $c \in [p]$. Let c' be the largest index such that $k^{c'} > 0$ and assume $c' > c$. Observe that, by hypothesis, $o^{c'} = 0$. The contribution of class c' and of the total congestion to $\psi(\boldsymbol{v})$ is $\overline{x}_{c'}k^{c'}(k^{<c'} + (k^{c'} + \delta)/2) - k(k+\delta)/2$. Let $\boldsymbol{v}'$ be the tuple obtained by removing all players of class c' from $\boldsymbol{v}$. The contribution of class c' and of the total congestion to $\psi(\boldsymbol{v}')$ is equal to $-(k - k^{c'})(k - k^{c'} + \delta)/2$.

So, as $\overline{x}_{c'} > 1$ by Lemma (4), we obtain

$$\begin{aligned}&\psi(\boldsymbol{v}) - \psi(\boldsymbol{v}') \\ &> k^{c'}\left(k^{<c'} + \frac{k^{c'} + \delta}{2}\right) - \frac{k(k+\delta)}{2} + (k - k^{c'})\frac{k - k^{c'} + \delta}{2} \\ &= k^{c'}(k^{<c'} + k^{c'} - k) \\ &= 0,\end{aligned}$$

where the last equality follows from the definition of c'. Thus, $\psi(\boldsymbol{v}') < \psi(\boldsymbol{v}) < 0$ implies that also $\psi(\boldsymbol{v}')$ is not satisfied by $\overline{\boldsymbol{s}}(p)$. By repeating the argument for all classes $c' > c$ with $k^{c'} > 0$, we get that there exists a constraint $\psi(\boldsymbol{v}')$ which is not satisfied by $\overline{\boldsymbol{s}}$ and such that $k^{c'} = 0$ for each $c' > c$. Now assume that $k^{c'} > 0$ for some $c' < c-1$. As the contribution of class c' to $\psi(\boldsymbol{v}')$ is strictly positive and $\overline{x}_{c'} > \overline{x}_{c-1}$ by claim *(i)* of Lemma 4, by transferring all the congestion $\sum_{c'<c-1} k^{c'}$ from all classes c' such that $c' < c-1$ to class $c-1$, we get a constraint which is not satisfied by $\overline{\boldsymbol{s}}$ and such that $k^{c'} = 0$ for each $c' < c-1$. Hence, we conclude that there exists a constraint which is almost homogeneous with respect to both $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma}^*$, has $o > 0$, and is not satisfied by $\overline{\boldsymbol{s}}(p)$: a contradiction.

Now, consider the case of $o = 0$. Again, as the contribution of class c to $\psi(\boldsymbol{v})$ is strictly positive for each $c \in [p]$ and, by claim *(i)* of Lemma 4, $\overline{x}_c > \overline{x}_p$ for each $c \in [p-1]$, by transferring all the congestion $\sum_{c<p} k^c$ from the first $p-1$ classes to class p, we get a constraint which is not satisfied by $\overline{\boldsymbol{s}}$ and such that $k^c = 0$ for each $c < p$. Hence, we conclude that there exists a constraint which is almost homogeneous with respect to both $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma}^*$, has $o = 0$, and is not satisfied by $\overline{\boldsymbol{s}}(p)$: a contradiction. □

By Lemma 7, to show that $\overline{\boldsymbol{s}}(p)$ is feasible for $DP(\Gamma_p)$, we need to prove that it satisfies every constraint belonging to each of the following five types:

t_0: $o = 0$ and $k = k^p > 0$,
t_1: $k = 0$ and $0 < o = o^c$ for some $c \in [p]$,
t_2: $0 < o = o^c$ and $0 < k = k^c$ for some $c \in [p]$,
t_3: $0 < o = o^c$ and $0 < k = k^{c-1}$ for some $c \in [p] \setminus \{1\}$,
t_4: $0 < o = o^c$, $0 < k = k^c + k^{c-1}$, and $k^c k^{c-1} \neq 0$ for some $c \in [p] \setminus \{1\}$.

We analyze each of them, separately, in the next subsections.

4.3.1. Constraints of type t_0

In this case, $o = 0$ and $k = k^p > 0$. So, the dual constraint becomes

$$\overline{x}_p \frac{k(k+\delta)}{2} - \frac{k(k+\delta)}{2} \geq 0$$

which is always satisfied as $\overline{x}_p > 1$ by Lemma 4.

4.3.2. Constraints of type t_1

In this case, $k = 0$ and $0 < o = o^c$ for some $c \in [p]$. So, the dual constraint becomes

$$-\overline{x}_c o\delta + \overline{\gamma}\frac{o(o+\delta)}{2} \geq 0,$$

which is always satisfied as $\overline{\gamma} \geq \overline{x}_1 \geq \overline{x}_c$ by constraint (7) of $NLP(p)$ and by claim *(i)* of Lemma 4 and since $o \geq \delta$ in both atomic and non-atomic games.

4.3.3. Constraints of type t_2

In this case, $0 < o = o^c$ and $0 < k = k^c$ for some $c \in [p]$. So, the dual constraint becomes

$$\overline{x}_c\left(\frac{k(k+\delta)}{2} - \frac{o(k+2\delta)}{2}\right) + \overline{\gamma}\frac{o(o+\delta)}{2} - \frac{k(k+\delta)}{2} \geq 0,$$

with $k, o > 0$ in general and $k, o \geq 1$ for atomic games in particular. By rearranging and setting $k = \theta o$, we get

$$\overline{\gamma} \geq \frac{\theta(\theta o + \delta) - \overline{x}_c(\delta(\theta-2) + o\theta(\theta-1))}{o+\delta}. \tag{13}$$

For non-atomic games, by substituting $\delta = 0$ in (13), we get $\overline{\gamma} \geq \theta(\theta - \overline{x}_c(\theta-1))$ whose right-hand side is maximized at $\theta^* = \frac{\overline{x}_c}{2(\overline{x}_c-1)}$ which yields $\overline{\gamma} \geq \frac{\overline{x}_c^2}{4(\overline{x}_c-1)}$. This inequality is satisfied for $c = p$ as $\overline{\gamma} \geq \frac{\overline{x}_{p-1}}{\overline{x}_{p-1}-1} = \frac{\overline{x}_p}{2(\overline{x}_p-1)} > \frac{\overline{x}_p^2}{4(\overline{x}_p-1)}$. Here, the last inequality comes from $\overline{x}_p < 2$. Moreover, $\overline{\gamma} \geq \frac{\overline{x}_c^2}{4(\overline{x}_c-1)}$ is also satisfied for any $c \in [p-1]$ as we have $\overline{x}_c \in [4/3, 4]$ by Lemma 4 and function $\frac{x^2}{4(x-1)}$ never exceeds the value $4/3$ for $x \in [4/3, 4]$.

For atomic games, denote by $g(o,\theta)$ the right-hand side of inequality (13) after substituting $\delta = 1$. As the sign of $\frac{\partial g}{\partial o}(o,\theta) = \frac{\theta(\theta-1)-\overline{x}_c(\theta^2-2\theta+2)}{(o+1)^2}$ is independent of o, it follows that $g(o,\theta)$ is maximized for either $o = 1$ or $o \to \infty$. We continue by analyzing these two cases.

Assume, first, that $o = 1$, which implies $\theta \in \mathbb{Z}$, with $\theta \geq 1$. We have $h(\theta) := g(1,\theta) = \frac{\theta(\theta+1)-\overline{x}_c(\theta^2-2)}{2}$. By computing $\frac{\partial h}{\partial \theta}$, we get that $h(\theta)$ is maximized at $\theta^* = \frac{1}{2(\overline{x}_c-1)}$. By claims *(ii)* and *(iii)* of Lemma 4, we know that $\overline{x}_c \geq 8/7$. This implies that

$\theta^* < 4$. Hence, as we also know that $\theta^* \in \mathbb{Z}$, with $\theta^* \geq 1$, we can claim that $h(\theta)$ is maximized for some $\theta^* \in [4]$, which gives $\overline{\gamma} \geq 1 + \overline{x}_c/2$, $\overline{\gamma} \geq 3 - \overline{x}_c$, $\overline{\gamma} \geq 6 - 7\overline{x}_c/2$ and $\overline{\gamma} \geq 10 - 7\overline{x}_c$. The first inequality is satisfied as $\overline{\gamma} \geq \overline{x}_1 = \overline{x}_1/2 + \overline{x}_1/2 \geq 1 + \overline{x}_c/2$. The second inequality is satisfied as $\overline{\gamma} \geq \overline{x}_1 \geq 2 > 3 - \overline{x}_c$. The third inequality is satisfied as $\overline{\gamma} \geq \overline{x}_1 \geq 2 \geq 6 - 7\overline{x}_p/2 \geq 6 - 7\overline{x}_c/2$. The last inequality is satisfied as $\overline{\gamma} \geq \overline{x}_1 \geq 2 \geq 10 - 7\overline{x}_p \geq 10 - 7\overline{x}_c$. All derivations follow from constraint (7) of $NLP(p)$ together with claims *(i)–(iv)* of Lemma 4.

Assume, now, that $o \to \infty$. Inequality (13) becomes $\overline{\gamma} \geq \theta(\theta - \overline{x}_c(\theta - 1))$, i.e., we obtain the same constraint already analyzed in the case of non-atomic games.

4.3.4. Constraints of type t_3

In this case, $0 < o = o^c$ and $0 < k = k^{c-1}$ for some $c \in [p] \setminus \{1\}$. So, the dual constraint becomes

$$\overline{x}_{c-1}\frac{k(k+\delta)}{2} - \overline{x}_c o\,(k+\delta) + \overline{\gamma}\frac{o(o+\delta)}{2} - \frac{k(k+\delta)}{2} \geq 0,$$

with $k, o > 0$ in general and $k, o \geq 1$ for atomic games in particular. By rearranging and setting $k = \theta o$, we get

$$\overline{\gamma} \geq \frac{(\theta o + \delta)(2\overline{x}_c - \theta\overline{x}_{c-1} + \theta)}{o + \delta}. \tag{14}$$

For non-atomic games, by substituting $\delta = 0$ in (14), we get $\overline{\gamma} \geq \theta(2\overline{x}_c - \theta\overline{x}_{c-1} + \theta)$ whose right-hand side is maximized at $\theta^* = \frac{\overline{x}_c}{\overline{x}_{c-1}-1}$ which yields $\overline{\gamma} \geq \frac{\overline{x}_c^2}{\overline{x}_{c-1}-1}$, that is, constraint (8) of $NLP(p)$.

For atomic games, denote by $g(o, \theta)$ the right-hand side of inequality (14) after substituting $\delta = 1$. As the sign of $\frac{\partial g}{\partial o}(o, \theta) = \frac{(\theta-1)^2(2\overline{x}_c - \theta(\overline{x}_{c-1}-1))}{(o+1)^2}$ is independent of o, it follows that $g(o, \theta)$ is maximized for either $o = 1$ or $o \to \infty$. We continue by analyzing these two cases.

Assume, first, that $o = 1$, which implies $\theta \in \mathbb{Z}$, with $\theta \geq 1$. We have $h(\theta) := g(1, \theta) = \frac{(\theta+1)(2\overline{x}_c - \theta\overline{x}_{c-1} + \theta)}{2}$. By computing $\frac{\partial h}{\partial \theta}$, we get that $h(\theta)$ is maximized at $\theta^* = \frac{1 + 2\overline{x}_c - \overline{x}_{c-1}}{2(\overline{x}_{c-1}-1)}$. The value of θ^* is increasing in $\overline{x}_c$, so, by claim *(i)* of Lemma 4, θ^* is maximized for $\overline{x}_c = \overline{x}_{c-1}$ which yields $\theta^* \leq \frac{\overline{x}_{c-1}+1}{2(\overline{x}_{c-1}-1)}$. Also, the value of θ^* is decreasing in $\overline{x}_{c-1}$, so, by claim *(ii)* of Lemma 4, θ^* is minimized for $\overline{x}_{c-1} = \overline{x}_c = 4/3$ which yields $\theta^* \geq 7/2$. Hence, as $\theta^* \in \mathbb{Z}$ and $\theta^* \geq 1$ by assumption, we can claim that $h(\theta)$ is maximized for some $\theta^* \in [4]$, which gives $\overline{\gamma} \geq 1 + 2\overline{x}_c - \overline{x}_{c-1}$, $\overline{\gamma} \geq 3 + 3\overline{x}_c - 3\overline{x}_{c-1}$, $\overline{\gamma} \geq 6 + 4\overline{x}_c - 6\overline{x}_{c-1}$ and $\overline{\gamma} \geq 10 + 5\overline{x}_c - 10\overline{x}_{c-1}$. All these four inequalities are implied by $\overline{\gamma} \geq \frac{\overline{x}_c^2}{\overline{x}_{c-1}-1}$ which comes from constraint (8) of $NLP(p)$.

Assume, now, that $o \to \infty$. Inequality (14) becomes $\overline{\gamma} \geq \theta(\theta + 2\overline{x}_c - \theta\overline{x}_{c-1})$, thus obtaining the same constraint analyzed in the case of non-atomic games.

4.3.5. Constraints of type t_4

In this case, $0 < o = o^c$, $0 < k = k^c + k^{c-1}$, and $k^c k^{c-1} \neq 0$ for some $c \in [p] \setminus \{1\}$. So, the dual constraint becomes

$$\overline{x}_{c-1}\frac{k^{c-1}(k^{c-1}+\delta)}{2} + \overline{x}_c\left(k^c\left(k^{c-1} + \frac{k^c+\delta}{2}\right) - o\left(k^{c-1} + \frac{k^c + 2\delta}{2}\right)\right) + \overline{\gamma}\frac{o(o+\delta)}{2} - \frac{(k^{c-1}+k^c)(k^{c-1}+k^c+\delta)}{2} \geq 0,$$

with $k^{c-1}, k^c, o > 0$ in general and $k^{c-1}, k^c, o \geq 1$ for atomic games in particular. By rearranging and setting $k^{c-1} = \theta o$ and $k^c = \psi o$, we get

$$\overline{\gamma} \geq \frac{-\theta\overline{x}_{c-1}(\theta o + \delta) - \overline{x}_c(\delta(\psi - 2) + o(\psi - 1)(2\theta + \psi))}{o + \delta} + \frac{(\theta + \psi)(\delta + (\theta + \psi)o)}{o + \delta}. \tag{15}$$

For non-atomic games, by substituting $\delta = 0$ in (15), we get $\overline{\gamma} \geq -\theta^2\overline{x}_{c-1} + (2\theta + \psi)(1 - \psi)\overline{x}_c + (\theta + \psi)^2$. Denote by $h(\theta, \psi)$ the right-hand side of this inequality. As the derivative $\frac{\partial h}{\partial \theta}(\theta, \psi) = -2(\theta(\overline{x}_{c-1} - 1) + \overline{x}_c(\psi - 1) - \psi)$ (resp. $\frac{\partial h}{\partial \psi}(\theta, \psi) = -2\psi(\overline{x}_c - 1) + \overline{x}_c(1 - 2\theta) + 2\theta$) is linear in θ (resp. ψ) and assumes a negative value for $\theta \to \infty$ (resp. $\psi \to \infty$), it follows that $h(\theta, \psi)$ is maximized at a value θ^* (resp. ψ^*) such that either $\theta^* = 0$ (resp. $\psi^* = 0$) or θ^* (resp. ψ^*) is such that $\frac{\partial h}{\partial \theta}(\theta^*, \psi) = 0$ (resp. $\frac{\partial h}{\partial \psi}(\theta, \psi^*) = 0$). In particular, the second option occurs if and only if $\theta^* > 0$ (resp. $\psi^* > 0$). Observe that, whenever $\theta^* = 0$ (resp. $\psi^* = 0$), a constraint of type t_4 boils down to a constraint of type t_2 (resp. type t_3) and we are done by the analysis provided in the previous subsections. So, we only need to consider the case of $\theta^* > 0$ and $\psi^* > 0$. By solving the system made of the two equations $\frac{\partial h}{\partial \theta}(\theta^*, \psi^*) = 0$ and $\frac{\partial h}{\partial \psi}(\theta^*, \psi^*) = 0$, we get $\theta^* = \frac{\overline{x}_c}{2(\overline{x}_{c-1} - \overline{x}_c)}$ and $\psi^* = \frac{\overline{x}_c(1 + \overline{x}_{c-1} - 2\overline{x}_c)}{2(\overline{x}_{c-1} - \overline{x}_c)(\overline{x}_c - 1)}$. Note that, by Lemma 5, conditions $\theta^* > 0$ and $\psi^* > 0$ require $c = p$. By substituting the values of θ^* and ψ^* in $h(\theta, \psi)$ and using $\overline{x}_p = \frac{2\overline{x}_{p-1}}{\overline{x}_{p-1}+1}$, we obtain constraint (9) of $NLP(p)$.

For atomic games, denote by $g(o,\theta,\psi)$ the right-hand side of inequality (15) after substituting $\delta=1$. As the sign of $\frac{\partial g}{\partial o}(o,\theta,\psi)=\frac{-\theta(\theta-1)\overline{x}_{c-1}+\overline{x}_c(\psi^2+2\psi(\theta-1)-2(\theta-1))-\psi^2+\psi(1-2\theta)-\theta(\theta-1)}{(o+1)^2}$ is independent of o, it follows that $g(o,\theta,\psi)$ is maximized for either $o=1$ or $o\to\infty$. We continue by analyzing these two cases.

Assume first that $o=1$, which implies $\theta\in\mathbb{Z}$, with $\theta\geq 1$. We have

$$\begin{aligned} h(\theta,\psi) &:= g(1,\theta,\psi) \\ &= \frac{-\theta(\theta+1)\overline{x}_{c-1}-(2\theta(\psi-1)+\psi^2-2)\overline{x}_c+(\theta+\psi)(\theta+\psi+1)}{2}. \end{aligned}$$

As the derivative $\frac{\partial h}{\partial\theta}$ (resp. $\frac{\partial h}{\partial\psi}$) is linear in θ (resp. ψ) and assumes a negative value for $\theta\to\infty$ (resp. $\psi\to\infty$), it follows that $h(\theta,\psi)$ is maximized at a value θ^* (resp. ψ^*) such that either $\theta^*=0$ (resp. $\psi^*=0$) or θ^* (resp. ψ^*) is such that $\frac{\partial h}{\partial\theta}(\theta^*)=0$ (resp. $\frac{\partial h}{\partial\psi}(\psi^*)=0$). In particular, the second option occurs if and only if $\theta^*>0$ (resp. $\psi^*>0$). Again, observe that, whenever $\theta^*=0$ (resp. $\psi^*=0$), a constraint of type t_4 boils down to a constraint of type t_2 (resp. type t_3) and we are done. So, we only need to consider the case of $\theta^*>0$ and $\psi^*>0$. By solving $\frac{\partial h}{\partial\theta}(\theta^*)=0$ and $\frac{\partial h}{\partial\psi}(\psi^*)=0$, we get $\theta^*=\frac{2\overline{x}_c-\overline{x}_{c-1}}{2(\overline{x}_{c-1}-\overline{x}_c)}$ and $\psi^*=\frac{\overline{x}_c(1+\overline{x}_{c-1}-2\overline{x}_c)}{2(\overline{x}_{c-1}-\overline{x}_c)(\overline{x}_c-1)}$. Note that, by Lemma 5, conditions $\theta^*>0$ and $\psi^*>0$ require $c=p$. By substituting the values of θ^* and ψ^* in $h(\theta,\psi)$ and using $\overline{x}_p=\frac{2\overline{x}_{p-1}}{\overline{x}_{p-1}+1}$ and $\overline{x}_{p-1}=\frac{\overline{\gamma}}{\overline{\gamma}-1}$, we obtain the inequality $\overline{\gamma}\geq\frac{8\overline{\gamma}^3-4\overline{\gamma}^2-2\overline{\gamma}-1}{8(\overline{\gamma}-1)(2\overline{\gamma}-1)}$ which is satisfied for any $\overline{\gamma}\geq 2$.

Assume now that $o\to\infty$. Inequality (15) becomes $\overline{\gamma}\geq-\theta^2\overline{x}_{c-1}+(2\theta+\psi)(1-\psi)\overline{x}_c+(\theta+\psi)^2$, thus obtaining the same constraint analyzed in the case of non-atomic games.

4.4. Upper bounds

Having shown that $\overline{\boldsymbol{s}}(p)$ is feasible for $DP(\Gamma_p)$, we can claim the following result.

Corollary 1. *For any priority-based affine congestion game Γ_p with $p\geq 2$, $\mathsf{PoA}(\Gamma_p)\leq\overline{\gamma}(p)$.*

By numerically solving $NLP(p)$, we explicitly quantify the upper bounds on the price of anarchy for some values of p as outlined in Table 1 (where, for completeness, we also report the previously known bound for the case of $p=1$, which is not covered by our analysis).

4.5. Lower bounds

Here, we construct, given an integer $p\geq 2$, a family of singleton congestion games to obtain lower bounds on the price of stability matching the upper bounds given in Corollary 1 for the price of anarchy. These games, which cover both the atomic and non-atomic cases, are defined by relying on the optimal solution $\overline{\boldsymbol{s}}(p)$ for $NLP(p)$. It is important to highlight that the explicit computation of $\overline{\boldsymbol{s}}(p)$ is not necessary. It is important to stress, here, that the structure of these lower bounding instances is obtained by implementing those constraints in $DP(\Gamma_p)$ giving life to the constraints defining $NLP(p)$, as illustrated in Observation 1.

Before presenting the promised family of games, we warm up by considering separately the cases of $p=2,3$ that require different constructions.

Theorem 4. *For any $\epsilon>0$, there exists a singleton atomic game Γ_2^a such that $\mathsf{PoS}(\Gamma_2^a)\geq 2-\epsilon$.*

Proof. Game Γ_2^a is defined as follows. There are θ players of class 1 and θ players of class 2. The set of resources R is defined as follows: $R=R_1\cup\{r_2\}$, with $R_1=\{r_{1,1},r_{1,2},\ldots,r_{1,\theta}\}$. All resources in R_1 have a linear latency function with coefficient equal to $(\theta+2)/2$, while resource r_2 has a linear latency function with coefficient equal to 1. All players of class 2 have a unique strategic choice[6] corresponding to resource r_2, while each player of class 1 can choose between two resources, called the *first* and *second* resource, respectively. More precisely, the ith player of class 1 can choose between resources $r_{1,i}$ and r_2. Observe that Γ_2^a is a singleton game.

It is immediate to check that the second strategy, which may cost at most $(\theta+1)/2$, is a dominant one for all players of class 1. Thus, the strategy profile $\boldsymbol{\sigma}$ in which all players of class 1 choose their second resource is the unique pure Nash equilibrium for Γ_2^a. We lower bound the price of stability of Γ_2^a by comparing the social cost of $\boldsymbol{\sigma}$ with the one yielded by the strategy profile $\boldsymbol{\sigma}^*$ in which all players of class 1 choose their first resource. In particular, we shall consider the limit of this lower bound for $\theta\to\infty$. We get

[6] This assumption is not unreasonable, as it is equivalent to the case in which all the alternatives available to the player have a latency function equal to $\ell(x)=\infty$, and so they are never chosen in a pure Nash equilibrium.

$$\lim_{\theta\to\infty} \mathsf{PoS}(\Gamma_2^a) \geq \lim_{\theta\to\infty} \frac{\mathsf{SC}(\boldsymbol{\sigma})}{\mathsf{SC}(\boldsymbol{\sigma}^*)} = \lim_{\theta\to\infty} \frac{\frac{1}{2}2\theta(2\theta+1)}{\theta\left(\frac{\theta+2}{2}\right)+\frac{1}{2}\theta(\theta+1)} = 2,$$

thus showing the claim. □

The previous construction can be easily adapted to provide a lower bound for the price of anarchy (and so also for the price of stability) of non-atomic games.

Theorem 5. *There exists a singleton non-atomic game Γ_2^{na} such that* $\mathsf{PoA}(\Gamma_2^{na}) \geq 2$.

Proof. Game Γ_2^{na} is defined as follows. Both classes have a flow of θ. The set of resources R is defined as follows: $R = \{r_1, r_2\}$. Resource r_1 has a constant latency function equal to $\theta/2$, while resource r_2 has a linear latency function with coefficient equal to 1. The flow of class 2 has a unique strategic choice corresponding to resource r_2, while the flow of class 1 can choose between both resources. Observe that Γ_2^{na} is a singleton game.

It is immediate to check that the strategy profile $\boldsymbol{\sigma}$ in which all flow of class 1 is assigned to r_2 is a pure Nash equilibrium for Γ_2^{na}. We lower bound the price of anarchy of Γ_2^{na} by comparing the social cost of $\boldsymbol{\sigma}$ with the one yielded by the strategy profile $\boldsymbol{\sigma}^*$ in which all flow of class 1 is assigned to r_1. We get

$$\mathsf{PoA}(\Gamma_2^{na}) \geq \frac{\mathsf{SC}(\boldsymbol{\sigma})}{\mathsf{SC}(\boldsymbol{\sigma}^*)} = \frac{\frac{(2\theta)^2}{2}}{\frac{\theta^2}{2}+\frac{\theta^2}{2}} = 2,$$

thus showing the claim. □

Theorem 6. *For any $\epsilon > 0$, there exists a singleton atomic game Γ_3^a such that* $\mathsf{PoS}(\Gamma_3^a) \geq 2.3247 - \epsilon$.

Proof. For a fixed value $x > 1$, Γ_3^a is defined as follows. There are $(x-1)^2\theta$ players of class 1, $(x-1)\theta$ players of class 2 and θ players of class 3. The set of resources R is defined as follows: $R = R_1 \cup \{r_2\} \cup \{r_3\}$, with $R_1 = \{r_{1,1}, r_{1,2}, \ldots, r_{1,(x-1)^2\theta}\}$. All resources in R_1 have a linear latency function with coefficient equal to $\frac{(x-1)\theta+3}{4}$, resource r_2 has a linear latency function with coefficient equal to $\frac{(x-1)\theta+2}{2(x-1)^2\theta+2}$, while resource r_3 has a linear latency function with coefficient equal to 1. All players of class 3 have a unique strategic choice corresponding to resource r_3, while, each player of class $c \in [2]$ can choose between two resources, called the *first* and *second* resource, respectively. More precisely, every player of class 2 can choose between r_2 and r_3, while the ith player of class 1 can choose between resources $r_{1,i}$ and r_2. Observe that Γ_3^a is a singleton game.

We show that the strategy profile $\boldsymbol{\sigma}$ in which all players of classes 1 and 2 choose their second resource is the unique pure Nash equilibrium for Γ_3^a. For any player of class 1, the first resource costs $\frac{(x-1)\theta+3}{4}$, while the second one may cost at most $\frac{(x-1)^2\theta+1}{2} \cdot \frac{(x-1)\theta+2}{2(x-1)^2\theta+2} = \frac{(x-1)\theta+2}{4}$. So, in any Nash equilibrium for Γ_3^a, all players of class 1 choose their second resource. Under this assumption, for any player of class 2, the first resource costs at least $((x-1)^2\theta+1)\frac{(x-1)\theta+2}{2(x-1)^2\theta+2} = \frac{(x-1)\theta+2}{2}$, while the second one may cost at most $\frac{(x-1)\theta+1}{2}$. Thus, $\boldsymbol{\sigma}$ is the unique pure Nash equilibrium for Γ_3^a.

We lower bound the price of stability of Γ_3^a by comparing the social cost of $\boldsymbol{\sigma}$ with the one yielded by the strategy profile $\boldsymbol{\sigma}^*$ in which all players of classes 1 and 2 choose their first resource. In particular, we shall consider the limit of this lower bound for $\theta \to \infty$. We get

$$\begin{aligned}\lim_{\theta\to\infty} \mathsf{PoS}(\Gamma_3^a) &\geq \lim_{\theta\to\infty} \frac{\mathsf{SC}(\boldsymbol{\sigma})}{\mathsf{SC}(\boldsymbol{\sigma}^*)}\\ &= \lim_{\theta\to\infty} \frac{\frac{(x-1)\theta+2}{2(x-1)^2\theta+2}\binom{(x-1)^2\theta+1)}{2}+\binom{\theta x+1)}{2}}{\frac{(x-1)\theta+3}{4}\binom{(x-1)^2\theta+1}{2}+\frac{(x-1)\theta+2}{2(x-1)^2\theta+2}\binom{(x-1)\theta+1}{2}+\binom{\theta+1}{2}}\\ &= \frac{x^3-x^2+3x-1}{x(x^2-3x+4)}.\end{aligned}$$

By choosing $x = 2.3247$, we get $\lim_{\theta\to\infty} \mathsf{PoS}(\Gamma_3^a) \geq 2.3247$ thus showing the claim. □

Again we easily extend the construction to deal with the price of anarchy of non-atomic games.

Theorem 7. *There exists a singleton non-atomic game Γ_3^{na} such that* $\mathsf{PoA}(\Gamma_3^{na}) \geq 2.3247$.

Proof. For a fixed value $x > 1$, game Γ_3^{na} is defined as follows. We have $f_1 = (x-1)^2\theta$, $f_2 = (x-1)\theta$ and $f_3 = \theta$. The set of resources R is defined as follows: $R = \{r_1, r_2, r_3\}$. Resource r_1 has a constant latency function equal to $(x-1)\theta/4$, resource r_2

has a linear latency function with coefficient equal to $\frac{1}{2(x-1)}$, while resource r_3 has a linear latency function with coefficient equal to 1. The flow of class 3 has a unique strategic choice corresponding to resource r_3, while each flow of class $i \in [2]$ can choose between resources r_i and r_{i+1}. Observe that Γ_3^{na} is a singleton game.

It is immediate to check that the strategy profile $\boldsymbol{\sigma}$ in which all flow of class $i \in [2]$ is assigned to r_{i+1} is a pure Nash equilibrium for Γ_3^{na}. We lower bound the price of anarchy of Γ_3^{na} by comparing the social cost of $\boldsymbol{\sigma}$ with the one yielded by the strategy profile $\boldsymbol{\sigma}^*$ in which all flow of class $i \in [2]$ is assigned to r_i. We get

$$\mathsf{PoA}(\Gamma_3^{na}) \geq \frac{\mathsf{SC}(\boldsymbol{\sigma})}{\mathsf{SC}(\boldsymbol{\sigma}^*)} = \frac{\frac{x^2\theta^2}{2} + \frac{(x-1)^4\theta^2}{4(x-1)}}{\frac{\theta^2}{2} + \frac{(x-1)^2\theta^2}{4(x-1)} + \frac{(x-1)^3\theta^2}{4}} = \frac{x^3 - x^2 + 3x - 1}{x(x^2 - 3x + 4)}.$$

Again, by choosing $x = 2.3247$, we get $\mathsf{PoA}(\Gamma_3^{na}) \geq 2.3247$ thus showing the claim. □

We now show how to generalize the previous constructions for any $p \geq 4$.

Theorem 8. *For any $\epsilon > 0$ and $p \geq 4$, there exists a priority-based singleton affine atomic congestion game Γ_p^a such that $\mathsf{PoS}(\Gamma_p^a) \geq \overline{\gamma}(p) - \epsilon$.*

Proof. Fix a value $\epsilon > 0$ and an integer $p \geq 4$ and consider the following singleton atomic game Γ_p^a. For every $c \in [p]$, the number of players of class c is equal to $|P_c| := \pi_c$, with

$$\pi_c = \begin{cases} \theta & \text{if } c = p, \\ \frac{\theta}{2(\overline{x}_{p-1}(p)-1)} & \text{if } c = p-1, \\ \frac{\overline{x}_{c+1}(p)}{\overline{x}_c(p)-1}\pi_{c+1} & \text{if } c \in [p-2]. \end{cases}$$

Here, the values $\overline{x}_c(p)$ for each $c \in [p]$ are the ones yielded by the optimal solution $\overline{\mathbf{s}}(p)$ for $NLP(p)$. We shall consider the case in which θ goes to infinity. Thus, as $\overline{x}_c(p) > 1$ for each $c \in [p]$ by Lemma 4, each value π_c belongs to the set of positive integers and is, so, well defined.

The set of resources is $R = R_1 \cup \{r_2\}, \ldots, \{r_p\}$, with $R_1 = \{r_{1,1}, \ldots, r_{1,|P_1|}\}$. All resources in R_1 have a linear latency function with coefficient equal to α_1, while, for $c \in [p] \setminus \{1\}$, resource r_c has a linear latency function with coefficient equal to α_c. All players of class p have a unique strategic choice corresponding to resource r_p. For each $c \in [p-1]$, instead, each player of class c can choose between two resources, called the *first* and *second* resource, respectively. For every $c \in [p-1] \setminus \{1\}$, the first and second resources of a player of class c are r_c and r_{c+1}, while the ith player of class 1 can choose between resources $r_{1,i}$ and r_2. Observe that Γ_p^a is a singleton game.

In order to maximize the price of anarchy yielded by this instance, let us use the pair of primal-dual formulations $PP(\Gamma_p^a)$ and $DP(\Gamma_p^a)$, where we set $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma}^*$ as the strategy profiles in which all players of class c, with $c \in [p-1]$, choose their second and first resource, respectively. As we consider the case of θ going to infinity, which implies that the number of players in each class grows arbitrarily large, we can get rid of small constants in the formulation, thus obtaining the following simplified primal linear program $PP(\Gamma_p^a)$:

$$\begin{aligned}
\max\ & \sum_{i=2}^{p-1} \frac{\alpha_i \pi_{i-1}^2}{2} + \frac{\alpha_p(\pi_{p-1}+\pi_p)^2}{2} \\
s.t.\ & \frac{\alpha_2\pi_1^2}{2} - \sum_{i\in[\pi_1]} \alpha_1 \leq 0, \\
& \frac{\alpha_{c+1}\pi_c^2}{2} - \alpha_c\pi_c\pi_{c-1} \leq 0, \quad \forall c \in [p-1] \setminus \{1\} \\
& \sum_{i\in[\pi_1]} \alpha_{1,i} + \sum_{i=2}^{p} \frac{\alpha_i \pi_i^2}{2} = 1 \\
& \alpha_c \geq 0, \quad \forall c \in [p]
\end{aligned}$$

The dual program $DP(\Gamma_p^a)$ is the following:

$$\begin{aligned}
\min\ & \gamma \\
s.t.\ & -x_1 + \gamma \geq 0 \\
& \frac{\pi_{c-1}^2}{2}x_{c-1} - \pi_{c-1}\pi_c x_c + \gamma\frac{\pi_c^2}{2} - \frac{\pi_{c-1}^2}{2} \geq 0 \qquad \forall c \in [p-1] \setminus \{1\}
\end{aligned}$$

$$\frac{\pi_{p-1}^2}{2}x_{p-1}+\gamma\frac{\pi_p^2}{2}-\frac{(\pi_{p-1}+\pi_p)^2}{2}\geq 0$$
$$x_c \geq 0 \qquad \forall c\in[p]$$

As $DP(\Gamma_p^a)$ is a particular instantiation of the dual program $DP(\cdot)$, it follows that its optimal solution has to be not smaller than $\overline{\gamma}(p)$. However, by substituting the values π_c and setting $x_c=\overline{x}_c(p)$ for each $c\in[p]$, $DP(\Gamma_p^a)$ rewrites as:

$$\begin{aligned}
&\min\ \gamma\\
&s.t.\ \overline{x}_1(p)\leq\gamma\\
&\quad \overline{x}_{c+1}(p)^2\leq\gamma(\overline{x}_c(p)-1) \qquad \forall c\in[p-2]\\
&\quad \overline{x}_{p-1}(p)\leq\gamma(\overline{x}_{p-1}(p)-1)\\
&\quad \overline{x}_p(p)=\frac{2\overline{x}_{p-1}(p)}{\overline{x}_{p-1}(p)+1}.
\end{aligned}$$

This implies that $\overline{\gamma}(p)$ is also an optimal solution for $DP(\Gamma_p^a)$.

Now consider the solution for $PP(\Gamma_p^a)$ obtained by setting $\alpha_p=1$ and $\alpha_c=\frac{\alpha_{c+1}\pi_c}{2\pi_{c-1}}$ for each $c\in[p-1]$, where we assume $P_0=\emptyset$, so that $\pi_0=0$. By the complementary slackness conditions, as this solution satisfies at equality all primal constraints which are related to a non-zero dual variable, it follows that $\mathsf{SC}(\boldsymbol{\sigma})/\mathsf{SC}(\boldsymbol{\sigma}^*)=\overline{\gamma}(p)$. This indeed shows that the price of anarchy of Γ_p^a is at least $\overline{\gamma}(p)$.

To extend this result to the price of stability, we need to show that $\boldsymbol{\sigma}$ is the unique pure Nash equilibrium for Γ_p^a. To this aim, we slightly perturb the coefficients of the latency functions by setting $\alpha_c=\frac{\alpha_{c+1}(\pi_c+1)}{2(\pi_{c-1}+1)}+\epsilon'$ for each $c\in[p-1]$, where $\epsilon'>0$ is arbitrarily small. With this modification, we prove that, for every $c\in[p-1]$, under the assumption that all players of class $c-1$ choose their second resource, playing the second resource is a dominant strategy for all players of class c. Because players of class $c-1$ are using their second resource, the first resource of a player of class c costs at least $(\pi_{c-1}+1)\alpha_c$, while the second one costs at most $\frac{(\pi_c+1)\alpha_{c+1}}{2}$. By the definition of α_c, the second resource always yields a strictly smaller cost, thus showing the claim. The modification decreases the ratio $\mathsf{SC}(\boldsymbol{\sigma})/\mathsf{SC}(\boldsymbol{\sigma}^*)$ of a negligible amount so that, for a suitable choice of ϵ', we have $\mathsf{PoS}(\Gamma_p^a)\geq\overline{\gamma}(p)-\epsilon$. □

The game used in the proof of the previous theorem can be adapted, with some modifications, to show the same result for non-atomic games.

Theorem 9. *For any $p\geq 4$, there exists a priority-based singleton affine non-atomic congestion game Γ_p^{na} such that $\mathsf{PoA}(\Gamma_p^{na})\geq\overline{\gamma}(p)$.*

Proof. Fix an integer $p\geq 4$ and consider the following singleton non-atomic game Γ_p^{na}. For every $c\in[p]$, the amount of flow f_c of class c is equal to

$$f_c=\begin{cases}1 & \text{if } c=p,\\ \frac{1}{2(\overline{x}_{p-1}(p)-1)} & \text{if } c=p-1,\\ \frac{\overline{x}_{c+1}(p)}{\overline{x}_c(p)-1}f_{c+1} & \text{if } c\in[p-2].\end{cases}$$

Here again, the values $\overline{x}_c(p)$ for each $c\in[p]$ are the ones yielded by the optimal solution $\overline{\boldsymbol{s}}(p)$ for $NLP(p)$. As $\overline{x}_c(p)>1$ for each $c\in[p]$ by Lemma 4, each value f_c belongs to the set of positive reals and is, so, well defined.

The set of resources is $R=\{r_1,\ldots,r_p\}$. Resource r_1 has a constant latency function with coefficient equal to β_1, while, for $c\in[p]\setminus\{1\}$, resource r_c has a linear latency function with coefficient equal to α_c. The flow of class p has a unique strategic choice corresponding to resource r_p. For each $c\in[p-1]$, instead, each flow of class c can choose between two resources, called the *first* and *second* resource, respectively. For every $c\in[p-1]$, the first and second resources of the flow of class c are r_c and r_{c+1}. Observe that Γ_p^{na} is a singleton game.

In order to maximize the price of anarchy yielded by this instance, we again use the pair of primal-dual formulations $PP(\Gamma_p^{na})$ and $DP(\Gamma_p^{na})$, where we set $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma}^*$ as the strategy profiles in which all flow of class c, with $c\in[p-1]$, is assigned to its second and first resource, respectively. We obtain the following primal linear program $PP(\Gamma_p^{na})$:

$$\max\ \sum_{i=2}^{p-1}\frac{\alpha_i f_{i-1}^2}{2}+\frac{\alpha_p(f_{p-1}+f_p)^2}{2}$$
$$s.t.$$
$$\frac{\alpha_2 f_1^2}{2}-\beta_1 f_1\leq 0,$$

$$\frac{\alpha_{c+1} f_c^2}{2} - \alpha_c f_{c-1} f_c \le 0, \quad \forall c \in [p-1] \setminus \{1\}$$
$$\beta_1 f_1 + \sum_{i=2}^{p} \frac{\alpha_i f_i^2}{2} = 1$$
$$\alpha_i \ge 0, \quad \forall i \in [p] \setminus \{1\}$$
$$\beta_1 \ge 0.$$

The dual program $DP(\Gamma_p^{na})$ is the following:

$$\begin{aligned}
&\min\ \gamma \\
&s.t. \\
&\quad -x_1 + \gamma \ge 0 \\
&\quad \frac{f_{c-1}^2}{2} x_{c-1} - f_{c-1} f_c x_c + \gamma \frac{f_c^2}{2} - \frac{f_{c-1}^2}{2} \ge 0 \qquad \forall c \in [p-1] \\
&\quad \frac{f_{p-1}^2}{2} x_{p-1} + \gamma \frac{f_p^2}{2} - \frac{(f_{p-1} + f_p)^2}{2} \ge 0 \\
&\quad x_c \ge 0 \qquad \forall c \in [p].
\end{aligned}$$

As $DP(\Gamma_p^{na})$ is a particular instantiation of the dual program $DP(\cdot)$, it follows that its optimal solution has to be not smaller than $\overline{\gamma}(p)$. However, by substituting the values f_c and setting $x_c = \overline{x}_c(p)$ for each $c \in [p]$, $DP(\Gamma_p^{na})$ rewrites as:

$$\begin{aligned}
&\min\ \gamma \\
&s.t. \\
&\quad \overline{x}_1(p) \le \gamma \\
&\quad \overline{x}_{c+1}(p)^2 \le \gamma(\overline{x}_c(p) - 1) \qquad \forall c \in [p-2] \\
&\quad \overline{x}_{p-1}(p) \le \gamma(\overline{x}_{p-1}(p) - 1) \\
&\quad \overline{x}_p(p) = \frac{2\overline{x}_{p-1}(p)}{\overline{x}_{p-1}(p) + 1}.
\end{aligned}$$

This implies that $\overline{\gamma}(p)$ is also an optimal solution for $DP(\Gamma_p^{na})$.

Now consider the solution for $PP(\Gamma_p^{na})$ obtained by setting $\alpha_p = 1$ and $\alpha_c = \frac{\alpha_{c+1} f_c}{2 f_{c-1}}$ for each $c \in [p-1]$, where we assume $P_0 = \emptyset$, so that $f_0 = 0$. By the complementary slackness conditions, as this solution satisfies at equality all primal constraints which are related to a non-zero dual variable, it follows that $\mathsf{SC}(\boldsymbol{\sigma})/\mathsf{SC}(\boldsymbol{\sigma}^*) = \overline{\gamma}(p)$. This indeed shows that the price of anarchy of Γ_p^{na} is at least $\overline{\gamma}(p)$. $\square$

5. Conclusions

We have given tight bounds for the price of anarchy and the price of stability of both atomic and non-atomic affine congestion games, under the assumption that the set of players is partitioned into $p \ge 2$ priority classes and the resources schedule their users according to a priority-based policy, breaking ties uniformly at random. These bounds hold even for load balancing games. Our findings outline an interesting separation between the case of $p \ge 2$ and the priority-free scenario of $p = 1$. The results are obtained by using the primal-dual method of [6]. An important consequence of this fact is that the upper bounds extend with no degradation to coarse correlated equilibria, as shown in [7].

There are several possible research directions that may be investigated. For instance, one can consider generalizations such as weighted players, polynomial latency functions, approximate Nash equilibria. Although the price of anarchy matches the price of stability even under singleton strategies, this may not be the case in presence of symmetric players or identical resources: both these restricted scenarios may hide useful properties. Moreover, as the lower bounding instances are based on a very constrained construction, it is also interesting to address special cases in which priority classes and strategies are restricted to obey particular relationships. An orthogonal approach may be that of considering the presence of a central authority which has the power of assigning priority classes to the players so as to induce games with low price of anarchy or price of stability.

Our randomized model assumes that players are risk neutral. Different behavior may arise under alternative models of risk averseness as investigated in [50] for the priority-free case.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Cosimo Vinci reports financial support was provided by INdAM-GNCS. Cosimo Vinci reports financial support was provided by Government of Italy Ministry of Education, University and Research.

Data availability

No data was used for the research described in the article.

Acknowledgements

We thank the reviewers of *Theoretical Computer Science* for their useful comments that helped to improve the presentation of this work.

Cosimo Vinci was partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR "Algorithms, Games, and Digital Markets" and INdAM-GNCS.

References

[1] S. Aland, D. Dumrauf, M. Gairing, B. Monien, F. Schoppmann, Exact price of anarchy for polynomial congestion games, SIAM J. Comput. 40 (5) (2011) 1211–1233.
[2] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, T. Roughgarden, The price of stability for network design with fair cost allocation, SIAM J. Comput. 38 (4) (2008) 1602–1623.
[3] B. Awerbuch, Y. Azar, L. Epstein, The price of routing unsplittable flow, in: Proc. of STOC, ACM, 2005, pp. 57–66.
[4] Y. Azar, L. Fleischer, K. Jain, V.S. Mirrokni, Z. Svitkina, Optimal coordination mechanisms for unrelated machine scheduling, Oper. Res. 63 (3) (2015) 489–500.
[5] K. Bhawalkar, M. Gairing, T. Roughgarden, Weighted congestion games: price of anarchy, universal worst-case examples, and tightness, ACM Trans. Econ. Comput. 2 (4) (2014) 1–23.
[6] V. Bilò, A unifying tool for bounding the quality of non-cooperative solutions in weighted congestion games, Theory Comput. Syst. 62 (5) (2018) 1288–1317.
[7] V. Bilò, On the robustness of the approximate price of anarchy in generalized congestion games, Theor. Comput. Sci. 906 (2022) 94–113.
[8] V. Bilò, A. Fanelli, M. Flammini, L. Moscardelli, Performances of one-round walks in linear congestion games, Theory Comput. Syst. 49 (1) (2011) 24–45.
[9] V. Bilò, C. Vinci, On the impact of singleton strategies in congestion games, in: Proc. of ESA, in: LIPIcs, vol. 87, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017, 17.
[10] V. Bilò, C. Vinci, Dynamic taxes for polynomial congestion games, ACM Trans. Econ. Comput. 7 (3) (2019) 15.
[11] V. Bilò, C. Vinci, On Stackelberg strategies in affine congestion games, Theory Comput. Syst. 63 (6) (2019) 1128–1249.
[12] V. Bilò, C. Vinci, Congestion games with priority-based scheduling, in: Proc. of SAGT, in: LNCS, vol. 12283, Springer, 2020, pp. 67–82.
[13] V. Bilò, L. Moscardelli, C. Vinci, Uniform mixed equilibria in network congestion games with link failures, in: Proc. of ICALP, in: LIPIcs, vol. 107, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018, 146.
[14] I. Caragiannis, Better bounds for online load balancing on unrelated machines, in: Proc. of SODA, ACM, 2008, pp. 972–981.
[15] I. Caragiannis, Efficient coordination mechanisms for unrelated machine scheduling, Algorithmica 66 (3) (2013) 512–540.
[16] I. Caragiannis, A. Fanelli, An almost ideal coordination mechanism for unrelated machine scheduling, in: Proc. of SAGT, in: LNCS, vol. 9928, Springer, 2016, pp. 315–326.
[17] I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, L. Moscardelli, Tight bounds for selfish and greedy load balancing, Algorithmica 61 (3) (2011) 606–637.
[18] I. Caragiannis, V. Gkatzelis, C. Vinci, Coordination mechanisms, cost-sharing, and approximation algorithms for scheduling, in: Proc. of WINE, in: LNCS, vol. 10660, Springer, 2017, pp. 74–87.
[19] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, Taxes for linear atomic congestion games, ACM Trans. Algorithms 7 (1) (2010) 13.
[20] G. Christodoulou, M. Gairing, Price of stability in polynomial congestion games, ACM Trans. Econ. Comput. 42 (2) (2016) 10.
[21] G. Christodoulou, M. Gairing, Y. Giannakopoulos, P.G. Spirakis, The price of stability of weighted congestion games, SIAM J. Comput. 48 (5) (2019) 1544–1582.
[22] G. Christodoulou, E. Koutsoupias, On the price of anarchy and stability of correlated equilibria of linear congestion games, in: Proc. of ESA, in: LNCS, vol. 3669, Springer, 2005, pp. 59–70.
[23] G. Christodoulou, E. Koutsoupias, The price of anarchy of finite congestion games, in: Proc. of STOC, ACM, 2005, pp. 67–73.
[24] G. Christodoulou, E. Koutsoupias, A. Nanavati, Coordination mechanisms, Theor. Comput. Sci. 410 (36) (2009) 3327–3336.
[25] G. Christodoulou, V.S. Mirrokni, A. Sidiropoulos, Convergence and approximation in potential games, Theor. Comput. Sci. 438 (2012) 13–27.
[26] J. Cohen, C. Dürr, N.K. Thang, Non-clairvoyant scheduling games, Theory Comput. Syst. 49 (1) (2011) 3–23.
[27] R. Cole, J.R. Correa, V. Gkatzelis, V.S. Mirrokni, N. Olver, Inner product spaces for minsum coordination mechanisms, in: Proc. of STOC, ACM, 2011, pp. 539–548.
[28] R. Cole, Y. Dodis, T. Roughgarden, How much can taxes help selfish routing?, J. Comput. Syst. Sci. 77 (3) (2006) 444–467.
[29] J.R. Correa, N.E. Stier-Moses, Stackelberg routing in atomic network games, Technical Report DRO-2007-03, Columbia Business School, 2007.
[30] P. von Falkenhausen, T. Harks, Optimal cost sharing for resource selection games, Math. Oper. Res. 38 (1) (2013) 184–208.
[31] B. Farzad, N. Olver, A. Vetta, A priority-based model of routing, Chic. J. Theor. Comput. Sci. (2008), 28 pp.
[32] D. Fotakis, Stackelberg strategies for atomic congestion games, Theory Comput. Syst. 47 (1) (2010) 218–249.
[33] D. Fotakis, P.G. Spirakis, Cost-balancing tolls for atomic network congestion games, Internet Math. 5 (4) (2008) 343–363.
[34] V. Gkatzelis, K. Kollias, T. Roughgarden, Optimal cost-sharing in general resource selection games, Oper. Res. 64 (6) (2016) 1230–1238.
[35] R. Gopalakrishnan, J.R. Marden, A. Wierman, Potential games are necessary to ensure pure Nash equilibria in cost sharing games, Math. Oper. Res. 39 (4) (2013) 1252–1296.
[36] T. Harks, I. Kleinert, M. Klimm, R.H. Möhring, Computing network tolls with support constraints, Networks 65 (3) (2015) 262–285.

[37] M. Hoefer, L. Olbrich, A. Skopalik, Taxing subnetworks, in: Proc. of WINE, in: LNCS, vol. 5385, Springer, 2008, pp. 286–294.
[38] S. Ieong, R. McGrew, E. Nudelman, Y. Shoham, Q. Sun, Fast and compact: a simple class of congestion games, in: Proc. of the 20th National Conference on Artificial Intelligence, AAAI, AAAI Press, 2005, pp. 489–494.
[39] N. Immorlica, L. Li, V.S. Mirrokni, A.S. Schulz, Coordination mechanisms for selfish scheduling, Theor. Comput. Sci. 410 (17) (2009) 1589–1598.
[40] T. Jelinek, M. Klaas, G. Schäfer, Computing optimal tolls with arc restrictions and heterogeneous players, in: Proc. of STACS, in: LIPIcs, vol. 25, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2014, pp. 433–444.
[41] G. Karakostas, S. Kolliopoulos, Stackelberg strategies for selfish routing in general multicommodity networks, Algorithmica 53 (1) (2009) 132–153.
[42] M. Klimm, D. Schmand, A. Tönnis, The online best reply algorithm for resource allocation problems, in: Proc. of SAGT, in: LNCS, vol. 11801, Springer, 2019, pp. 200–215.
[43] E. Koutsoupias, C. Papadimitriou, Worst-case equilibria, in: Proc. of STACS, in: LNCS, vol. 1653, Springer, 1999, pp. 404–413.
[44] D. Monderer, L.S. Shapley, Potential games, Games Econ. Behav. 14 (1996) 124–143.
[45] J.F. Nash, Equilibrium points in n-person games, Proc. Natl. Acad. Sci. 36 (1) (1950) 48–49.
[46] N. Nisan, T. Roughgarden, Vazirani V.V. Tardos, Algorithmic Game Theory, Cambridge University Press, 2007.
[47] D. Paccagnan, R. Chandan, B.L. Ferguson, J.R. Marden, Optimal taxes in atomic congestion games, ACM Trans. Econ. Comput. 9 (3) (2021) 19.
[48] D. Paccagnan, M. Gairing, In congestion games, taxes achieve optimal approximation, in: Proc. of EC, ACM Press, 2021, pp. 743–744.
[49] A.C. Pigou, The Economics of Welfare, Macmillan, 1920.
[50] G. Piliouras, E. Nikolova, J.S. Shamma, Risk sensitivity of price of anarchy under uncertainty, ACM Trans. Econ. Comput. 5 (1) (2016) 5.
[51] R.W. Rosenthal, A class of games possessing pure-strategy Nash equilibria, Int. J. Game Theory 2 (1) (1973) 65–67.
[52] T. Roughgarden, Stackelberg scheduling strategies, SIAM J. Comput. 33 (2) (2004) 332–350.
[53] T. Roughgarden, Intrinsic robustness of the price of anarchy, J. ACM 62 (5) (2015) 32.
[54] T. Roughgarden, E. Tardos, How bad is selfish routing?, J. ACM 49 (2) (2002) 236–259.
[55] S. Suri, C. Tóth, Y. Zhou, Selfish load balancing and atomic congestion games, Algorithmica 47 (1) (2007) 79–96.
[56] C. Swamy, The effectiveness of Stackelberg strategies and tolls for network congestion games, ACM Trans. Algorithms 8 (4) (2012) 36.
[57] C. Vinci, Non-atomic one-round walks in congestion games, Theor. Comput. Sci. 764 (2019) 61–79.