# Minimizing Rosenthal's Potential in Monotone Congestion Games

Vittorio Bilò
University of Salento
Lecce, Italy
vittorio.bilo@unisalento.it

Angelo Fanelli
Université Paris-Dauphine, Université PSL, CNRS, LAMSADE
Paris, France
angelo.fanelli@dauphine.psl.eu

Laurent Gourvès
Université Paris-Dauphine, Université PSL, CNRS, LAMSADE
Paris, France
laurent.gourves@lamsade.dauphine.fr

Christos Tsoufis
Université Paris-Dauphine, Université PSL, CNRS, LAMSADE
Paris, France
christos.tsoufis@dauphine.eu

Cosimo Vinci
University of Salento
Lecce, Italy
cosimo.vinci@unisalento.it

## ABSTRACT

Congestion games are attractive because they can model many concrete situations where some competing entities interact through the use of some shared resources, and also because they always admit pure Nash equilibria which correspond to the local minima of a potential function. We explore the problem of computing a state of minimum potential in this setting. Using the maximum number of resources that a player can use at a time, and the possible symmetry in the players' strategy spaces, we settle the complexity of the problem for instances having monotone (i.e., either non-decreasing or non-increasing) latency functions on their resources. The picture, delineating polynomial and NP-hard cases, is complemented with tight approximation algorithms.

## KEYWORDS

Congestion Games; Potential Function; Pure Equilibrium; Approximation

**ACM Reference Format:**
Vittorio Bilò, Angelo Fanelli, Laurent Gourvès, Christos Tsoufis, and Cosimo Vinci. 2025. Minimizing Rosenthal's Potential in Monotone Congestion Games. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 9 pages.

## 1 INTRODUCTION

Congestion games form one of the most studied classes of games in (Algorithmic) Game Theory. They provide a model of competition among $n$ strategic players requiring the use of certain resources in a set of $m$ available ones. Every resource has a cost function, also called latency function in the realm of transportation and routing networks, which only depends on the number of its users (a.k.a. the resource congestion). Given a state of the game in which all players have performed a strategic choice, the cost of a player is defined by the sum of the costs of all the selected resources.



*Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA*. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

Congestion games owe their success to the fact that they can model several concrete scenarios of competition, such as traffic networks, scheduling, group formation and cost sharing, to name a few [36]. At the same time, they possess intriguing and useful theoretical properties. In fact, Rosenthal [30] proved that every congestion game admits an exact potential (*Rosenthal's potential*): a function defined from the set of states of the game to the reals such that, every time a player performs a deviation from a state to another, the difference in the potential equals the difference of the player's costs in the two states. For finite games, this implies that every sequence of deviations in which a player improves her cost must have finite length and end at a pure Nash equilibrium, which is a state in which no player can improve her cost by changing her strategic choice. Years later, Monderer and Shapley [26] complemented this result by showing that every game admitting an exact potential is isomorphic to a congestion game.

Several algorithmic questions pertaining congestion games and their notable variants have been posed and addressed in the literature. Among these are computing a Nash equilibrium [18], computing a state minimizing the sum of the players' costs (a.k.a. the *social optimum*) [25, 28], bounding the worst-case (price of anarchy [23]) and the best-case (price of stability [3]) approximation of the social optimum yielded by a pure Nash equilibrium, and computing a state minimizing Rosenthal's potential [16, 18, 22].

The latter problem, in particular, has interesting applications. First, by definition, every local minimum of an exact potential function corresponds to a pure Nash equilibrium in the game. Therefore, computing the global minimum of Rosenthal's potential directly provides a pure Nash equilibrium for the given congestion game. Additionally, (approximate) potential minimizers often yield good approximations of the social optimum, offering pure Nash equilibria whose efficiency can approach or even match the (approximate) price of stability [11, 14]. Finally, when minimizing the potential function is intractable, approximating the potential becomes crucial: By finding states with low potential, better-response dynamics can lead to equilibria with similarly low potential, providing an efficient path to desirable outcomes with minimal social cost.

Fabrikant *et al.* [18] were the first to attack this problem. They show how to compute the minimum of Rosenthal's potential in symmetric network congestion games with non-decreasing latency functions through a reduction to a min-cost flow problem. In network congestion games, resources are edges in a graph and every player wants to select a path connecting a source to a destination; it is symmetric when all players share the same source-destination pair. Ackermann *et al.* [1] extended this result to the case in which all players share the same source (or, equivalently, the same destination) only. Del Pia *et al.* [16] and Kleer and Schäfer [22] adopt a polyhedral approach to solve the problem under certain structural properties of the players' strategic space, still in the case of non-decreasing latency functions. They assume that the incidence vectors of the strategies of a player are given by the binary vectors in a certain polytope. Del Pia *et al.* [16], in particular, provide a solution for *symmetric totally unimodular congestion games*, i.e., for the case in which the matrix defining the common polytope encoding the strategies of all players is totally unimodular. Kleer and Schäfer [22] further generalize the result to the cases in which the polytope obtained by aggregating the polytopes encoding the strategies of each player satisfies two properties named, respectively, *integer decomposition property* (IDP) and *box-totally dual integrality property* (box-TDI).

Apart from its connections to the search of pure Nash equilibria, another interesting application of the potential minimization problem is related to optimal job scheduling, and in particular, arises in group-based variants of load balancing with related machines and identical jobs, where the jobs are initially partitioned into $n$ groups, and each group $i \in [n]$ consists of $m_i$ jobs. Each group $i$ of $m_i$ jobs can be assigned to a set $S$ of $m_i$ distinct machines (assigning each job to a different machine), chosen from a collection of feasible sets of $m_i$ machines. For example, consider the scenario where machines are positioned within a metric space, and a set $S$ of $m_i$ machines is deemed feasible if all machines in $S$ are within a certain distance from one another (for instance, to enable cooperation within the same group). The objective is to find an assignment of jobs to machines that minimizes the total (or average) completion time. If the completion time of any job assigned to a machine depends on the number of jobs allocated to it, we observe that minimizing the total completion time of the jobs can be reduced to minimizing the potential function.[1]

In this work, we continue the study of the problem of computing a state minimizing Rosenthal's potential, that we refer to as Min Potential, and depart from previous approaches in what follows. First, rather than considering the combinatorial structure of the players' strategy space, we look at the maximum number of resources that a player can use simultaneously. Secondly, besides of the case of non-decreasing latency functions, which is a typical assumption in road and communication networks where congestion has a detrimental effect on the cost of using a resource, we also consider non-increasing functions, which is typical in cost-sharing scenarios [3].

## 1.1 Our Contribution

For games with non-decreasing latency functions, we obtain a precise characterization of Min Potential with respect to the maximum number of resources that a player can use simultaneously (a.k.a. the *size*). The results, which also depend on whether players' strategy spaces are symmetric or not, are summarized in Table 1.

**Table 1: Games with non-decreasing latency functions: Summary of the complexity results with respect to both the size and the symmetry (symmetric & general) of the players' strategy space**

| | size = 1 | size = 2 | size ≥ 3 |
|---|---|---|---|
| sym. | $O(nm)$ (Cor. 1) | $O(n^3m^3)$ (Th. 2) | NP-hard (Th. 4) |
| gen. | $O(n^3m^3)$ ([22]) | NP-hard (Th. 3) | NP-hard (Th. 4) |

Given the hardness results stated in Theorems 3 and 4 (see Table 1), we also focus on the computation of good approximate solutions to Min Potential. We heavily exploit an approximation algorithm designed by Paccagnan and Gairing [28] for the problem of computing the social optimum in congestion games with non-negative, non-decreasing and semi-convex latency functions. The approximation guarantee, which depends in a fairly complicated way on the values of these functions, is proved to be tight, unless P = NP. For polynomial latency functions of maximum degree $d$ (and non-negative coefficients), Paccagnan and Gairing show that the bound simplifies to the $(d+1)$-th Bell number, denoted as $B_{d+1}$. This result uses (a generalization of) Dobinski's formula [24].

We show how their algorithm can be used to provide approximate solutions to Min Potential as well. This is done by observing that Min Potential on a congestion game with non-negative and non-decreasing latency functions can be reduced to the problem of computing a social optimum on the same game with perturbed latency functions which are non-negative, non-decreasing and semi-convex. So, Paccagnan and Gairing's algorithm can be applied. However, the resulting approximation factor is expressed as an infinite sum that, in general, does not have an explicit representation. Thus, the results of Paccagnan and Gairing cannot be directly applied to quantify the approximation factor for Min Potential in a simple and explicit way. A possible approach to address this problem could be to apply Dobinski's formula to obtain a simpler bound. Indeed, in the case of polynomial latency functions of maximum degree $d$, the reduction produces a game whose latency functions are still polynomials of maximum degree $d$. Nevertheless, the resulting polynomials are quite specific and may have negative coefficients. Therefore, Dobinski's formula cannot be directly applied to derive tight and explicit bounds on the approximation guarantee, or to show that $B_{d+1}$ continues to hold at least as an upper bound.

As our major contribution, we provide a highly non-trivial analysis of the above approximation guarantee, by which we derive a precise bound equal to $\Lambda_d := \sum_{j=0}^{d} \frac{j+2}{j+1} {d \brace j}$, where ${d \brace j}$ denotes the Stirling numbers of the second kind (see Section 2 for formal definitions). It is easy to check that, for any $d \geq 1$, $\Lambda_d$ never exceeds $\frac{3}{2}B_d$, with the inequality being tight only for the case of $d = 1$, and that this value is always smaller than $B_{d+1}$. Moreover, given that $B_d$ grows asymptotically as $(f(d))^d$ with $f(d) = \Theta(d/\ln(d))$ [5], it follows that the difference between $B_{d+1}$ and $\Lambda_d$ increases with $d$. A comparison between $B_{d+1}$ and $\Lambda_d$ for small values of $d$ is

[1]Group-based variants of job scheduling, which are similar but not equivalent to ours, have also been studied (e.g., in [37]).

reported in Table 2. Last but not least, since the inapproximability result provided by Paccagnan and Gairing holds for any class of latency functions, we immediately obtain that the approximation guarantee of $\Lambda_d$ is tight for Min Potential.

**Table 2: Comparison between the tight approximation guarantee for the problem of minimizing the social cost (equal to $B_{d+1}$) and for Min Potential (equal to $\Lambda_d$), when considering polynomial latency functions of maximum degree $d$**

| | $d=1$ | $d=2$ | $d=3$ | $d=4$ | $d=5$ | $d=6$ | $d=7$ |
|---|---|---|---|---|---|---|---|
| $B_{d+1}$ | 2 | 5 | 15 | 52 | 203 | 877 | 4140 |
| $\Lambda_d$ | 1.5 | 2.84 | 6.75 | 19.54 | 65.92 | 251.98 | 1070.21 |

For games with non-increasing latency functions, Min Potential shows to be generally harder, see Table 3. In fact, while a solution can be easily computed in the case of symmetric games of constant size, the problem becomes NP-hard as soon as we drop the symmetry assumption, and this holds even if $\texttt{size} = 1$, all resources share the same latency function, and players only have two possible strategies. For general latencies, $\texttt{size} = 1$, and no specific limit on the number of possible strategies for the players, we show that Min Potential cannot be approximated to better than $H_n = \Theta(\ln n)$, unless P = NP, and we provide a matching approximation guarantee (Theorem 11).

**Table 3: Summary of the complexity results for constant size games with non-increasing latency functions**

| | $\texttt{size} = O(1)$ |
|---|---|
| sym. | $m^{O(1)}$ (Prop. 3) |
| gen. | NP-hard when $\texttt{size} = 1$ (Th. 10) |

For all missing material of this article, we refer the reader to [6].

## 1.2 Further Related Work

The problem of computing a global minimum of Rosenthal's potential is a specialization of that of computing a local minimum for this function. This problem, which is equivalent to computing a pure Nash equilibrium for a given game, has received quite some attention in the literature of congestion games. However, while for non-decreasing latency functions a series of results [1, 13, 15, 17–19, 21] has provided a fairly complete understanding of the complexity of this problem, much less in known for the case of non-increasing latency functions [2, 10, 33].

Our approximation for Min Potential with polynomial latency functions is obtained by exploiting an algorithm designed by Paccagnan and Gairing [28]. This algorithm uses taxes to force selfish uncoordinated players to implement provably efficient solutions. The efficiency of taxation mechanisms in congestion games with non-decreasing latency functions has been studied in a series of papers [7–9, 12, 27, 28, 34, 35]. In [7–9, 12, 27, 28, 35], the aim is to use taxes to lead selfish agents towards states with provably good social cost, while, in [34], the authors also consider the objective of minimizing the *stretch*: a worst-case measure of the discrepancy between the potential of a pure Nash equilibrium and the optimal social cost. This measure has application in the computation of approximate pure Nash equilibria.

# 2 PRELIMINARIES

## 2.1 Mathematical Definitions

Given a positive integer $k$, we denote by $[k]$ the set $\{1, 2, \dots, k\}$. Given two integers $d$ and $k$ with $0 \le k \le d$, the *Stirling number of the second kind*, denoted $\left\{{d \atop k}\right\}$, is the number of ways to partition a set of $d$ elements into $k$ non-empty subsets. As such, they obey the following recursive definition: $\left\{{d+1 \atop k}\right\} = k\left\{{d \atop k}\right\} + \left\{{d \atop k-1}\right\}$. Some simple identities involving these numbers that hold essentially by definition are: $\left\{{d \atop d}\right\} = 1$, $\left\{{d \atop 1}\right\} = 1$ for every $d \ge 1$, and $\left\{{d \atop 2}\right\} = 2^{d-1} - 1$. It has been proved, see [29], that $\left\{{d \atop k}\right\} \ge \frac{k^2+k+2}{2} k^{d-k-1} - 1$. Using this lower bound, it is immediate to show that $\left\{{d \atop d-2}\right\} \ge \frac{d^3-5d^2+10d-10}{2}$, with the right-hand side always increasing in $d$, which implies $\left\{{d \atop d-2}\right\} \ge 7$ for every $d \ge 4$. For a given $d \ge 0$, the *Bell number*, denoted $B_d$, counts the number of possible partitions of a set of $d$ elements. By definition, it immediately follows that $B_d = \sum_{k=0}^{d} \left\{{d \atop k}\right\}$.[2] For two non-negative integers $i$ and $j$, the *falling factorial*, denoted $(i)_j$, is defined as $(i)_j := i \cdot (i-1) \dots \cdot (i-j+1) = \prod_{k=0}^{j-1}(i-k)$, with the convention that $(i)_j := 0$ for $j = 0$.

## 2.2 The Model

A congestion game $\mathcal{G}$ is a tuple $\langle N, R, (S_i)_{i\in N}, (\ell_r)_{r\in R}\rangle$. $N$ denotes the set of *players* and $R$ the set of *resources*. We assume that both $N$ and $R$ are finite and non-empty and define $n := |N|$ and $m := |R|$. Each player $i \in N$ is associated with a finite and non-empty set of *strategies* $S_i \subseteq 2^R$. If every strategy in $S_i$ consists of one resource then we say that $\mathcal{G}$ is a *singleton* game. If all players share the same set of strategies, i.e., $S_i = S_j$ for every $i, j \in N$, then we say that $\mathcal{G}$ is a *symmetric* game (in that case, $S$ denotes the strategy set of all players). We denote by $\texttt{size}(\mathcal{G})$ the maximum cardinality of any strategy, i.e., $\texttt{size}(\mathcal{G}) := \max_{i\in N} \max_{s \in S_i} |s|$. Hence, a singleton game $\mathcal{G}$ is such that $\texttt{size}(\mathcal{G}) = 1$. Every resource $r \in R$ is associated with a *latency function* $\ell_r : \mathbb{N} \mapsto \mathbb{R}_{\ge 0}$, which maps the number of users of $r$ to a non-negative real. We assume that $\ell_r$ is monotone for all $r \in R$; we also suppose that $\ell_r(0) = 0$ and $\ell_r(1) > 0$. Function $\ell_r$ is *non-decreasing* (resp., *non-increasing*) when $\ell_r(h) \ge \ell_r(h')$ (resp., $\ell_r(h) \le \ell_r(h')$) for every $h > h' \ge 1$. Sections 3 and 4 deal with instances where every latency function is non-decreasing, whereas Section 5 is devoted to instances where every latency function is non-increasing. We say that $\ell_r$ is *polynomial* of maximum degree $d \in \mathbb{N}$ if $\ell_r(x) = \sum_{q=0}^{d} \alpha_{r,q} x^q$, for some coefficients $\alpha_{r,0}, \dots, \alpha_{r,d} \ge 0$; it is *affine* if it is polynomial of maximum degree 1 and is *linear* if it is affine and $\alpha_{r,0} = 0$.

The set of *states* of the game is denoted by $\mathbf{S} := S_1 \times S_2 \times \dots \times S_n$. The $i$-th component of a state $\mathbf{s} \in S$ is the strategy played by player $i$ in $\mathbf{s}$ and is denoted by $\mathbf{s}_i$. For every state $\mathbf{s}$ and resource $r$, we

[2] Given that $\left\{{d \atop 0}\right\} = 0$ for any $d > 0$, this identity can be rewritten as $B_d = \sum_{k=1}^{d} \left\{{d \atop k}\right\}$, whenever $d > 0$.

denote by $n_r(\mathbf{s})$ the number of players using resource $r$ in $\mathbf{s}$, i.e., $n_r(\mathbf{s}) := |\{i \in N : r \in \mathbf{s}_i\}|$, and we refer to it as the *congestion* of $r$ in $\mathbf{s}$. For every state $\mathbf{s}$, the *cost* incurred by player $i$ in $\mathbf{s}$ is $c_i(\mathbf{s}) := \sum_{r \in \mathbf{s}_i} \ell_r(n_r(\mathbf{s}))$. Notice that, by definition of latency, $c_i(\mathbf{s}) > 0$ for every player $i$ and state $\mathbf{s}$.

## 2.3 Improving Moves, Potential Function and Pure Nash Equilibria

Let us consider a congestion game $\mathcal{G} = \langle N, R, (S_i)_{i \in N}, (\ell_r)_{r \in R} \rangle$. For every state $\mathbf{s} \in \mathrm{S}$, every player $i \in N$ and every $s \in S_i$, we denote by $[\mathbf{s}_{-i}, s]$ the new state obtained from $\mathbf{s}$ by setting the $i$-th component, that is the strategy of $i$, to $s$ and keeping all the remaining components unchanged, i.e., if $\bar{\mathbf{s}} = [\mathbf{s}_{-i}, s]$ then $\bar{\mathbf{s}}_i = s$ and $\bar{\mathbf{s}}_j = \mathbf{s}_j$ for every player $j \neq i$.

A congestion game is a strategic game in which every player $i$ selects $s \in S_i$ so as to minimize $c_i([\mathbf{s}_{-i}, s])$. The transition from $\mathbf{s}$ to $[\mathbf{s}_{-i}, s]$ is called a *move* of player $i$ from state $\mathbf{s}$. We say that a transition from $\mathbf{s}$ to $[\mathbf{s}_{-i}, s]$ is an *improving move* for $i$ if $c_i([\mathbf{s}_{-i}, s]) < c_i(\mathbf{s})$. We say that a state-valued function $\Gamma : \mathrm{S} \mapsto \mathbb{R}_{\geq 0}$ is an *exact potential function* for the game if $\Gamma(\mathbf{s}) - \Gamma([\mathbf{s}_{-i}, s]) = c_i(\mathbf{s}) - c_i([\mathbf{s}_{-i}, s])$ holds for every $\mathbf{s} \in \mathrm{S}$ and $s \in S_i$. This means that, in games admitting an exact potential $\Gamma$, if a player $i$ can make a move from $\mathbf{s}$ to $[\mathbf{s}_{-i}, s]$ such that $\Gamma(\mathbf{s}) > \Gamma([\mathbf{s}_{-i}, s])$, then the move must be improving for $i$, and the decrease in cost $c_i(\mathbf{s}) - c_i([\mathbf{s}_{-i}, s])$ for player $i$ is exactly $\Gamma(\mathbf{s}) - \Gamma([\mathbf{s}_{-i}, s])$. Meanwhile, the existence of an improving move from $\mathbf{s}$ to $[\mathbf{s}_{-i}, s]$ by player $i$ implies that $\Gamma(\mathbf{s}) > \Gamma([\mathbf{s}_{-i}, s])$, and the decrease in potential $\Gamma(\mathbf{s}) - \Gamma([\mathbf{s}_{-i}, s])$ is precisely $c_i(\mathbf{s}) - c_i([\mathbf{s}_{-i}, s])$. Therefore, the number of states $|\mathrm{S}|$ being finite, every maximal sequence of improvement moves leads to a *local minimum* of $\Gamma$, i.e., to a state in which no further improvement move can be performed.

Such a state is called *pure Nash equilibrium*. In other words, we say that a state $\mathbf{s} \in \mathrm{S}$ is a *pure Nash equilibrium* if, for every player $i \in N$ and every strategy $s \in S_i$, we have $c_i(\mathbf{s}) \leq c_i([\mathbf{s}_{-i}, s])$. It is well known that $\Phi_{\mathcal{G}}(\mathbf{s}) := \sum_{r \in R} \sum_{j=0}^{n_r(\mathbf{s})} \ell_r(j)$, called the Rosenthal's potential function [30], is an exact potential function for $\mathcal{G}$. Notice that, by definition of latency, $\Phi_{\mathcal{G}}(\mathbf{s}) > 0$ holds for every state $\mathbf{s}$.

## 2.4 Problem Statement

In this work, we are interested in the following problem, that we name Min Potential: given a congestion game $\mathcal{G}$, find a state of $\mathcal{G}$ minimizing $\Phi_{\mathcal{G}}$. Another interesting problem in the realm of congestion games is Min Social Cost, which, given a congestion game $\mathcal{G}$, asks for a state minimizing the *social cost* $\mathrm{SC}_{\mathcal{G}}(\mathbf{s}) := \sum_{i \in N} c_i(\mathbf{s})$ of $\mathcal{G}$, i.e., the sum of the costs of all players.

Given $\rho \geq 1$, a $\rho$-approximate state $\mathbf{s}$ is a feasible state of $\mathcal{G}$ for which $\mathrm{f}(\mathbf{s}) \leq \rho\, \mathrm{f}(\mathbf{s}^*)$ where $\mathbf{s}^*$ is a global minimizer of function f, where f can be either $\Phi_{\mathcal{G}}$ or $\mathrm{SC}_{\mathcal{G}}$. A $\rho$-approximation algorithm is a polynomial time algorithm which always outputs a $\rho$-approximate state.

# 3 COMPLEXITY OF MIN POTENTIAL WITH NON-DECREASING LATENCIES

In this section and the following one, we assume that every latency function is non-decreasing.

We start by considering the complexity of Min Potential for the basic case of $\mathtt{size}(\mathcal{G}) = 1$, i.e., the case of singleton congestion games. It is well known that any singleton congestion game can be interpreted also as a network one. Thus, the algorithm by Fabrikant *et al.* [18] for symmetric network congestion games can be applied to symmetric singleton congestion games as well. The reduction of Fabrikant *et al.* to min-cost flow produces a parallel-link graph with $nm$ edges. Given that the best algorithm for min-cost flow in a graph with $\alpha$ nodes and $\beta$ edges has complexity $O(\alpha\beta \log \alpha(\beta + \alpha \log \alpha))$ (see Armstrong and Jin [4]), it follows that Min Potential can be solved in $O(n^2 m^2)$ using approaches from the current state of the art.

We give a better algorithm exploiting the fact that, in singleton games, any sequence of improving moves has polynomially bounded length. Together with next proposition, showing that, if the game is symmetric, any local minimum of Rosenthal's potential is also a global minimum, it yields an $O(nm)$ algorithm.

**Proposition 1.** All pure Nash equilibria of a symmetric congestion game $\mathcal{G}$ with $\mathtt{size}(\mathcal{G}) = 1$ have the same potential.

Proof. Let $\mathbf{e}^*$ be a state with minimum potential. Clearly $\mathbf{e}^*$ is a pure Nash equilibrium. Assume by contradiction that there exists another pure Nash equilibrium $\mathbf{e}$ such that $\Phi_{\mathcal{G}}(\mathbf{e}^*) < \Phi_{\mathcal{G}}(\mathbf{e})$. Let us denote by $C(\mathbf{e})$ the set of all pure Nash equilibria obtained from $\mathbf{e}$ by renaming the players, i.e., $C(\mathbf{e}) = \{\mathbf{s} \in \mathrm{S} : n_r(\mathbf{s}) = n_r(\mathbf{e}) \text{ for all } r \in R\}$. Observe that all states in $C(\mathbf{e})$ are equilibria and have the same potential. For any state $\mathbf{s} \in \mathrm{S}$, let us denote by $\mathtt{over}(\mathbf{s}) \subseteq R$ the set of resources whose congestion in $\mathbf{s}$ is strictly larger than the congestion in $\mathbf{e}$ (or any other equilibrium in $C(\mathbf{e})$) and by $\mathtt{under}(\mathbf{s}) \subseteq R$ the set of resources whose congestion in $\mathbf{s}$ is strictly smaller than the congestion in $\mathbf{e}$, i.e., $\mathtt{over}(\mathbf{s}) = \{r \in R : n_r(\mathbf{s}) > n_r(\mathbf{e})\}$ and $\mathtt{under}(\mathbf{s}) = \{r \in R : n_r(\mathbf{s}) < n_r(\mathbf{e})\}$. Notice that, as long as $\mathbf{s}$ does not belong to $C(\mathbf{e})$, then both $\mathtt{over}(\mathbf{s})$ and $\mathtt{under}(\mathbf{s})$ are non-empty.

Let us consider a sequence $\mathbf{e}^* = \mathbf{s}^0, \mathbf{s}^1, \ldots, \mathbf{s}^k = \mathbf{e}'$ of $k+1 \geq 2$ states in which $\mathbf{e}' \in C(\mathbf{e})$ and, for every $t \in [0, \ldots, k-1]$, $\mathbf{s}^{t+1}$ is obtained from $\mathbf{s}^t$ by a move of player $\pi(t)$ who deviates from a resource in $\mathtt{over}(\mathbf{s}^t)$ to a resource in $\mathtt{under}(\mathbf{s}^t)$, i.e., $\mathbf{s}^t_{\pi(t)} \in \mathtt{over}(\mathbf{s}^t)$ and $\mathbf{s}^{t+1}_{\pi(t)} \in \mathtt{under}(\mathbf{s}^t)$. Notice that such sequence always exists, and that every move in the sequence decreases $\sum_{r \in R} |n_r(\mathbf{e}') - n_r(\mathbf{s}^t)|$ which is a measure of distance between the congestion vector of any member of $C(\mathbf{e})$ and the congestion vector of $\mathbf{s}^t$, i.e., $\sum_{r \in R} |n_r(\mathbf{e}') - n_r(\mathbf{s}^t)| > \sum_{r \in R} |n_r(\mathbf{e}') - n_r(\mathbf{s}^{t+1})|$ for all $t \in [0, \ldots, k-1]$. Since $\Phi_{\mathcal{G}}(\mathbf{e}^*) < \Phi_{\mathcal{G}}(\mathbf{e}')$, there must exist a time $t$ such that $\Phi_{\mathcal{G}}(\mathbf{s}^{t+1}) > \Phi_{\mathcal{G}}(\mathbf{s}^t)$; let $h < k$ be the first of such time steps, i.e., $\Phi_{\mathcal{G}}(\mathbf{s}^0) = \Phi_{\mathcal{G}}(\mathbf{s}^1) = \ldots = \Phi_{\mathcal{G}}(\mathbf{s}^h) < \Phi_{\mathcal{G}}(\mathbf{s}^{h+1})$. Let us assume that $\pi(h)$ at time $h$ is moving from resource $r$ to $r'$, i.e., $r = \mathbf{s}^h_{\pi(h)}$ and $r' = \mathbf{s}^{h+1}_{\pi(h)}$. By assumption, $\mathbf{s}^h$ is a state with minimum potential and therefore an equilibrium, while $\mathbf{s}^{h+1}$ is not an equilibrium – in fact the move of player $\pi(h)$ from resource $r'$ to $r$, leading from state $\mathbf{s}^{h+1}$ to $\mathbf{s}^h$, decreases the potential and hence is an improving move. Moreover, observe that, since $r \in \mathtt{over}(\mathbf{s}^h)$ and $r' \in \mathtt{under}(\mathbf{s}^h)$, we have that $n_r(\mathbf{s}^{h+1}) = n_r(\mathbf{s}^h) - 1 \geq n_r(\mathbf{e}')$ and $n_{r'}(\mathbf{s}^{h+1}) = n_{r'}(\mathbf{s}^h) + 1 \leq n_{r'}(\mathbf{e}')$. Combining these latter observations with the fact that the latencies are non-decreasing and that the move of $\pi(h)$ from $r'$ to $r$ in $\mathbf{s}^{h+1}$

is an improving move, we get that also the move of any player in $\mathbf{e}'$ from resource $r'$ to $r$ is an improving move, which implies that $\mathbf{e}'$ is not an equilibrium. Hence, a contradiction. □

**Corollary 1.** For every symmetric congestion game $\mathcal{G}$ with $\texttt{size}(\mathcal{G}) = 1$, Min Potential can be solved in $O(nm)$ time.

Proof. Given Proposition 1, it follows that any pure Nash equilibrium for a symmetric singleton congestion game is a solution to Min Potential. A Nash equilibrium in this setting can be easily computed as follows. Start from the empty state and let players sequentially choose the cheapest resource, given the choices of her predecessors. Every single player's decision can be done in $O(m)$ time, for a total complexity of $O(nm)$. The outcome is a pure Nash equilibrium because every player's decision is a best response with respect to her predecessors. This is also true for the successors. Indeed, if a successor of player $i$, say $j$, plays the same resource as $i$, then $i$ cannot profitably deviate because both $i$ and $j$ have the same strategy space. If $j$ plays a different resource from $i$, then $i$ cannot profitably move towards $j$'s resource because the latency functions are non-decreasing. □

Proposition 1 extends neither to non-symmetric singleton games nor to symmetric games of size two (see [6]).

Now, let us shift towards singleton games when the symmetry property is dropped. This case is also in P, since it is covered by the results of Kleer and Schäfer on Min Potential for polytopal congestion games satisfying some structural properties (namely, IDP and box-TDI) [22, Theorem 3.3]. An alternative (and more direct) way to prove this is to reduce the problem to finding a minimum weight perfect matching in a bipartite graph $(V_1 \cup V_2, E)$ such that $V_1 = N \cup D$, with $D$ being a set of $(m-1) \cdot n$ dummy vertices, $V_2 = \{(r, \mu) : r \in R \text{ and } \mu \in [n]\}$, and there is an edge of weight $\ell_r(\mu)$ between $i \in V_1$ and $(r, \mu) \in V_2$ if, and only if, $r \in S_i$, and there is an edge of weight 0 for all pair $(d, (r, \mu)) \in D \times V_2$. Thus, the problem can be solved in $O(n^3 m^3)$.

We now move on to the case of $\texttt{size}(\mathcal{G}) = 2$ and show that symmetry makes a huge difference here, as it creates a separation between tractable and intractable cases.

**Theorem 2.** For every symmetric congestion game $\mathcal{G}$ with $\texttt{size}(\mathcal{G}) = 2$, Min Potential can be solved in $O(n^3 m^3)$.

Proof sketch. We exploit a reduction to the problem of computing a Maximum Weight Matching of a given size. The input is a graph $G = (V, E)$, a weight function $w : E \longrightarrow \mathbb{R}_{\geq 0}$ and a positive integer $q$ such that $G$ admits a matching of size $q$. The problem is to find a matching $M \subseteq E$ of size exactly $q$ which maximizes $w(M) = \sum_{e \in M} w(e)$. The problem is known to be polynomial time solvable.[3]

Take a symmetric congestion game $\mathcal{G} = \langle N, R, S, (\ell_r)_{r \in R} \rangle$ with $\texttt{size}(\mathcal{G}) = 2$, where $S$ denotes the strategy space of all players. We can suppose without loss of generality that every strategy in $S$ consists of exactly two resources. To do so, we introduce a fictitious resource $r_0$ (namely, $R \longleftarrow R \cup \{r_0\}$) so that every singleton strategy $\{r_j\} \in S$ is replaced by $\{r_0, r_j\}$.

Now, we can construct an instance $I$ of the matching problem as follows. For every resource $r_i \in R$, we build a set of exactly $n$ vertices $\{v_i^1, \ldots, v_i^n\}$. Next, for every pair $\{r_j, r_k\} \in S$, we construct a complete bipartite graph between $\{v_j^1, \ldots, v_j^n\}$ and $\{v_k^1, \ldots, v_k^n\}$. Every edge $(v_j^a, v_k^b)$, where $j, k \neq 0$, has weight equal to $C - \ell_{r_j}(a) - \ell_{r_k}(b)$, where $C \geq 2 \cdot \max_{r \in R} \ell_r(n)$ and every edge $(v_j^a, v_0^b)$, where $j \neq 0$, has weight equal to $C - \ell_{r_j}(a)$.

**Claim 1.** A state $\mathbf{s}$ in $\mathcal{G}$ gives a matching $M$ of size $n$ in $I$ with weight $w(M) = n \cdot C - \Phi_{\mathcal{G}}(\mathbf{s})$.

**Claim 2.** A matching $M$ of size $n$ in $I$ gives a state $\mathbf{s}$ in $\mathcal{G}$ with potential $\Phi_{\mathcal{G}}(\mathbf{s}) \leq n \cdot C - w(M)$.

Now, the technique is to compute a Maximum Weight Matching $M$ of size $n$ in $I$. Consequently, from Claim 2 we have a state $\mathbf{s}$ with potential $\Phi_{\mathcal{G}}(\mathbf{s}) \leq n \cdot C - w(M)$. Next, assume that $\mathcal{G}$ admits a state $\mathbf{s}^*$ such that $\Phi_{\mathcal{G}}(\mathbf{s}^*) < \Phi_{\mathcal{G}}(\mathbf{s})$. Then, from Claim 1 we get a matching $M^*$ of size $n$ with weight $w(M^*) = n \cdot C - \Phi_{\mathcal{G}}(\mathbf{s}^*)$. However, using the hypothesis, we get that $w(M^*) > n \cdot C - \Phi_{\mathcal{G}}(\mathbf{s}) \geq w(M)$, which is a contradiction.

Concerning time complexity, computing a maximum weight matching of given size reduces to computing a maximum weight matching in a modified graph whose number of vertices is at most doubled. Computing a maximum weight matching is cubic in the number of vertices. Our initial graph having $nm$ vertices, the time complexity of our procedure is $O(n^3 m^3)$. □

When the symmetry property is dropped, Min Potential becomes intractable when $\texttt{size}(\mathcal{G}) = 2$.

**Theorem 3.** Min Potential is NP-hard for congestion games $\mathcal{G}$ with $\texttt{size}(\mathcal{G}) = 2$ and linear latencies.

Finally, we show that hardness of computation extends to even symmetric games as soon as $\texttt{size}(\mathcal{G})$ gets equal to three.

**Theorem 4.** Min Potential is NP-hard for symmetric congestion games $\mathcal{G}$ with $\texttt{size}(\mathcal{G}) \geq 3$ and linear latencies.

Since, for $\texttt{size}(\mathcal{G}) \geq 3$, Min Potential is NP-hard for the symmetric case, it is also NP-hard for the general case. We observe that, if one modifies the latency functions used in the proofs of Theorems 3 and 4 to be such that $\ell_r(1) = 1$ and $\ell_r(h) = M_\rho$ for each $h \geq 2$, where $M_\rho$ is an appropriate large number, then no $\rho$-approximation algorithm for Min Potential can be proposed, unless P = NP.

## 4 APPROXIMATING MIN POTENTIAL WITH POLYNOMIAL LATENCIES

In this section, we show how to achieve an optimal approximation for Min Potential, when considering general congestion games with polynomial latency functions of maximum degree $d \in \mathbb{N}$.

[3] See [32, Chapter 18.5f] for bipartite graphs. The case of non-bipartite graphs can be reduced to the traditional maximum weight matching problem by modifying the instance as follows: increase the weight of every edge by a positive constant $Z$ such that every matching of size $k$ has larger weight than any other matching of size $k-1$. Then, add some extra $|V| - 2q$ dummy vertices, and link them with the original vertices with edges whose weight $W$ is big enough so that any maximum weight matching $\mathcal{M}$ in the new graph must include $|V| - 2q$ edges saturating all the dummy vertices. Apart from these $|V| - 2q$ heavy edges, $2q$ vertices remain to be saturated, which is done with a matching $M \subset \mathcal{M}$ of cardinality $q$ whose edges all belong to the initial graph, and $M$ has maximum weight in the initial graph.

To show our result, we exploit an optimal approximation algorithm developed by Paccagnan and Gairing [28] for Min Social Cost. This algorithm applies to congestion games with very general latency functions (satisfying mild assumptions only), and the resulting approximation factor is represented as an infinite sum that depends on the values of these functions. Then, by exhibiting the equivalence between Min Social Cost and Min Potential, we show how to convert the approximation guarantee obtained by Paccagnan and Gairing for Min Social Cost into an optimal approximation for Min Potential. Then, we specialize the result to the case of polynomial latency functions of maximum degree $d$ and, by exploiting some techniques arising from combinatorics, we achieve an exact quantification of the approximation factor in terms of a weighted finite sum of Stirling numbers of the second kind. In particular, we will show that, for any fixed $\epsilon > 0$, Min Potential admits a $(\Lambda_d + \epsilon)$-approximation, with

$$\Lambda_d := \sum_{j=1}^{d} \left(\frac{j+2}{j+1}\right) \left\{ {d \atop j} \right\} \in \left[B_d, \frac{3}{2}B_d\right]. \tag{1}$$

We point out that most of the algorithmic machinery used to obtain the desired approximation relies on the work of Paccagnan and Gairing, and our careful analysis specializes their results to Min Potential applied to games with polynomial latency functions. Some values of $\Lambda_d$ are provided in Table 2.

## 4.1 Approximation Algorithm for Min Social Cost: a Quick Overview

For a given $y \in \mathbb{N}$ and a latency function $\tilde{\ell}$, let[4]

$$\rho_{\tilde{\ell}}(y) := \sup_{y\in\mathbb{N}} \frac{\mathbb{E}_{P\sim \mathrm{Poi}(y)}[P\tilde{\ell}(P)]}{y\tilde{\ell}(y)} = \frac{\sum_{x=0}^{\infty} x\tilde{\ell}(x)\frac{y^x}{x!e^y}}{y\tilde{\ell}(y)}, \tag{2}$$

with $\mathrm{Poi}(y)$ denoting the Poisson distribution with parameter $y$; furthermore, for a given class of latency functions $\tilde{\mathcal{L}}$, define

$$\rho_{\tilde{\mathcal{L}}} := \sup_{\tilde{\ell}\in\tilde{\mathcal{L}}} \sup_{y\in\mathbb{N}} \rho_{\tilde{\ell}}(y). \tag{3}$$

For a given (and arbitrarily small) $\epsilon > 0$ and a class of latency functions $\tilde{\mathcal{L}}$ which are non-negative, non-decreasing and semi-convex (i.e., such that function $\tilde{g}(x) = x\tilde{\ell}(x)$ is convex for any $\tilde{\ell} \in \tilde{\mathcal{L}}$), the approximation algorithm provided by Paccagnan and Gairing, denoted here as AlgMinCost, guarantees a $(\rho_{\tilde{\mathcal{L}}} + \epsilon)$-approximation to Min Social Cost, when applied to a congestion game $\tilde{\mathcal{G}}$ with latency functions in $\tilde{\mathcal{L}}$. Paccagnan and Gairing also show that the obtained approximation is essentially optimal. Indeed, they show that it is NP-hard to approximate Min Social Cost within a factor better than $\rho_{\tilde{\mathcal{L}}}$, when restricting to congestion games with latencies in $\tilde{\mathcal{L}}$, for any class of latency functions $\tilde{\mathcal{L}}$.

## 4.2 Min Potential versus Min Social Cost

Given a latency function $\ell$, let $\tilde{\ell}$ denote the latency function defined as $\ell(x) = \sum_{h=1}^{x} \ell(h)/x$ and, given a class of latency functions $\mathcal{L}$, let $\tilde{\mathcal{L}} := \{\tilde{\ell} : \ell \in \mathcal{L}\}$; analogously, given a congestion games $\mathcal{G}$ with latency functions $(\ell_r)_{r\in R}$, let $\tilde{\mathcal{G}}$ be the congestion game obtained from $\mathcal{G}$ by replacing each latency $\ell_r$ with $\tilde{\ell}_r$.

By resorting to the following proposition, that shows the equivalence between Min Social Cost and Min Potential, we will see how to apply AlgMinCost to Min Potential in order to have the same approximation guaranteed for Min Social Cost but on a narrowed set of latency functions.

**Proposition 2.** Let $\mathcal{G}$ be a congestion game with latency functions $(\ell_r)_{r\in R}$. Then: (i) the latency functions of $\tilde{\mathcal{G}}$ are non-negative, non-decreasing and semi-convex; (ii) the potential function $\Phi_{\mathcal{G}}$ of $\mathcal{G}$ coincides with the social cost $\mathrm{SC}_{\tilde{\mathcal{G}}}$ of $\tilde{\mathcal{G}}$.

By combining Proposition 2 with the findings of Paccagnan and Gairing, we obtain in polynomial time a $(\rho_{\tilde{\mathcal{L}}} + \epsilon)$-approximate solution for Min Potential as follows: starting from the input congestion game $\mathcal{G}$, we first construct the corresponding congestion game $\tilde{\mathcal{G}}$ (this can be done in polynomial time); then, by applying AlgMinCost we obtain a $(\rho_{\tilde{\mathcal{L}}} + \epsilon)$-approximate solution $\mathbf{s}$ for game $\tilde{\mathcal{G}}$, w.r.t. Min Social Cost; finally, as the potential function $\Phi_{\mathcal{G}}$ of $\mathcal{G}$ and the social cost $\mathrm{SC}_{\tilde{\mathcal{G}}}$ of $\tilde{\mathcal{G}}$ have the same value for all states (by Proposition 2), we have that $\mathbf{s}$ is also a $(\rho_{\tilde{\mathcal{L}}}+\epsilon)$-approximate solution for game $\mathcal{G}$, w.r.t. Min Potential. Furthermore, by the hardness results of Paccagnan and Gairing and the above observations, we have that the obtained approximation is essentially optimal for Min Potential (up to the arbitrarily small constant $\epsilon$).

## 4.3 Characterization of the Approximation Factor for Polynomial Latencies

Let $\mathcal{L}_d$ denote the class of polynomial latency functions of maximum degree $d$. By the above observations, AlgMinCost can be adapted to return a $(\rho_{\tilde{\mathcal{L}}_d} + \epsilon)$-approximation to Min Potential. However $\rho_{\tilde{\mathcal{L}}_d}$, as it is represented in definition (3) (reported from [28]), is defined in terms of an infinite sum, whose exact value can only be approximated and does not allow for a direct quantification of the asymptotic growth of the approximation factor as a function of $d$. In the following, we show that $\rho_{\tilde{\mathcal{L}}_d}$ coincides with the value $\Lambda_d$ defined in (1), and this characterization leads to a simpler and more precise estimation of the approximation ratio. A similar result has been obtained for the Min Social Cost problem by Paccagnan and Gairing [28], who showed, by exploiting the Dobinski's formula [24], that their tight approximation factor coincides with $B_{d+1}$. However, it seems that their analysis cannot be directly applied to Min Potential to obtain the same or similar bounds (see [6]).

**Theorem 5.** For any $d \in \mathbb{N}$, we have $\rho_{\tilde{\mathcal{L}}_d} = \Lambda_d \sim B_d \leq \left(\frac{0.792d}{\ln(d+1)}\right)^d$.

Proof. Fix $d \in \mathbb{N}$. We first observe that $\Lambda_d \sim B_d$ holds by the right-hand part of (1) and $B_d \leq \left(\frac{0.792d}{\ln(d+1)}\right)^d$ has been shown in [5]. Then, in the remainder of the proof we will focus on equality $\rho_{\tilde{\mathcal{L}}_d} = \Lambda_d$.

Let $\ell_d$ denote the monomial latency function defined as $\ell_d(x) = x^d$. We will first show that $\rho_{\{\tilde{\ell}_d\}} = \Lambda_d$, that is, we are restricting the original class of latency function $\mathcal{L}_d$ to the simple monomial function of degree $d$; then, we will generalize this restricted claim

[4]Since in this section we are going to consider a reduction from Min Potential on a game $\mathcal{G}$ to Min Social Cost on a game $\tilde{\mathcal{G}}$, we shall add a "tilde" to the notation pertaining the Min Social Cost problem on a game $\tilde{\mathcal{G}}$.

to the whole class $\mathcal{L}_d$ of polynomial latency functions of maximum degree $d$. For any $y \in \mathbb{N}$, define

$$\Lambda_d(y) := \frac{\sum_{j\in[d]} \left\{ {d \atop j} \right\} \left(\frac{y^{j+1}}{j+1} + y^j\right)}{\sum_{j\in[d]} \left\{ {d \atop j} \right\} \frac{(y+1)_{j+1}}{j+1}}. \tag{4}$$

By exploiting definition (2), we have

$$\rho_{\tilde{\ell}_d} = \frac{\sum_{x=0}^{\infty} x\tilde{\ell}_d(x)\frac{y^x}{x!e^y}}{y\tilde{\ell}_d(y)} = \frac{\sum_{x=0}^{\infty}\sum_{h\in[x]} h^d \frac{y^x}{x!e^y}}{\sum_{h\in[y]} h^d}. \tag{5}$$

The following lemma provides an alternative representation for the sum of the first $y$ $d$-th powers, in terms of Stirling numbers of the second kind.

**Lemma 6.** For any $y \in \mathbb{N}$, we have

$$\sum_{h\in[y]} h^d = \sum_{j\in[d]} \left\{ {d \atop j} \right\} \frac{(y+1)_{j+1}}{j+1}.$$

The equality reported in the above lemma is folklore and the proof can be found, for instance, in [20]. The following lemma, together with Lemma 6, will be used to show that (4) and (5) are distinct representations for the same number.

**Lemma 7.** For any $y \in \mathbb{N}$, we have that

$$\sum_{x=0}^{\infty}\sum_{h\in[x]} h^d \frac{y^x}{x!e^y} = \sum_{j\in[d]} \left\{ {d \atop j} \right\} \left(\frac{y^{j+1}}{j+1} + y^j\right).$$

Proof of Lemma 7. By applying Lemma 6 with $x$ in place of $y$, we have that

$$\begin{aligned}\sum_{x=0}^{\infty}\left(\sum_{h\in[x]} h^d\right)\frac{y^x}{x!e^y} &= \sum_{x=0}^{\infty}\left(\sum_{j\in[d]} \left\{ {d \atop j} \right\} \frac{(x+1)_{j+1}}{j+1}\right)\frac{y^x}{x!e^y} \\ &= \sum_{j\in[d]} \frac{1}{(j+1)e^y} \left\{ {d \atop j} \right\} \sum_{x=0}^{\infty} (x+1)_{j+1}\frac{y^x}{x!}. \end{aligned} \tag{6}$$

For any $j, y \in \mathbb{N}$, we have

$$\begin{aligned}\sum_{x=0}^{\infty}(x+1)_{j+1}\frac{y^x}{x!} &= y^j \sum_{x=0}^{\infty}\left(\frac{\partial}{\partial y}\right)^{j+1}\left(y\cdot\frac{y^x}{x!}\right) \\ &= y^j \left(\frac{\partial}{\partial y}\right)^{j+1}\left(y\sum_{x=0}^{\infty}\left(\frac{y^x}{x!}\right)\right) \\ &= y^j \left(\frac{\partial}{\partial y}\right)^{j+1}\left(ye^y\right) = y^j\left(ye^y + (j+1)e^y\right) \\ &= (j+1)e^y\left(\frac{y^{j+1}}{j+1} + y^j\right), \end{aligned} \tag{7}$$

where $\left(\frac{\partial}{\partial y}\right)^k (g(y))$ denotes the $k$-th derivative of $g$, the second and the third equality hold since the series of functions $\sum_{x=0}^{\infty} f_x$, with $f_x(t) = t\left(\frac{t^x}{x!}\right)$ for any $t \geq 0$ and $x \in \mathbb{N}$, uniformly converges to function $f(t) = te^t$ on non-negative closed intervals, and then the series of the derivatives converges to the derivative of the series (this last property is folklore and a proof is given, for instance, in [31]). Finally, by applying (7) to (6), we get

$$\begin{aligned}\sum_{x=0}^{\infty}\sum_{h\in[x]} h^d \frac{y^x}{x!e^y} &= \sum_{j\in[d]} \frac{1}{(j+1)e^y}\left\{ {d \atop j} \right\}\sum_{x=0}^{\infty}(x+1)_{j+1}\frac{y^x}{x!} \\ &= \sum_{j\in[d]} \left\{ {d \atop j} \right\}\left(\frac{y^{j+1}}{j+1} + y^j\right),\end{aligned}$$

that shows the claim. □

By Lemma 6 and Lemma 7 we get

$$\begin{aligned}\rho_{\tilde{\ell}_d}(y) &= \frac{\sum_{x=0}^{\infty}\sum_{h\in[x]} h^d \frac{y^x}{x!e^y}}{\sum_{h\in[y]} h^d} = \frac{\sum_{x=0}^{\infty}\sum_{h\in[x]} h^d \frac{y^x}{x!e^y}}{\sum_{j\in[d]} \left\{ {d \atop j} \right\} \frac{(y+1)_{j+1}}{j+1}} \\ &= \frac{\sum_{j\in[d]} \left\{ {d \atop j} \right\} \left(\frac{y^{j+1}}{j+1} + y^j\right)}{\sum_{j\in[d]} \left\{ {d \atop j} \right\} \frac{(y+1)_{j+1}}{j+1}} = \Lambda_d(y), \end{aligned} \tag{8}$$

for any $y \in \mathbb{N}$.

The following lemma shows that, independently of the value of $d$, the maximum over $y$ of $\Lambda_d(y)$ is given by $y = 1$.

**Lemma 8.** For any $y \in \mathbb{N}$, we have $\Lambda_d(1) \geq \Lambda_d(y)$.

By putting the above results together, we get

$$\rho_{\{\tilde{\ell}_d\}} = \sup_{y\in\mathbb{N}} \rho_{\tilde{\ell}_d}(y) = \sup_{y\in\mathbb{N}} \Lambda_d(y) = \Lambda_d(1), \tag{9}$$

where the second and the third equality hold by equality (8) and Lemma 8, respectively. We observe that (9) shows the claim of the theorem, when restricting the class $\mathcal{L}_d$ to the monomial function $\ell_d(x) = x^d$ only. We will generalize (9) to the whole class $\mathcal{L}_d$ of polynomial latency functions of maximum degree $d$. To do this, it is sufficient to show that

$$\rho_{\tilde{\ell}^*_d}(y) \leq \Lambda_d(1), \tag{10}$$

for any $\ell^*_d \in \mathcal{L}_d$ and $y \in \mathbb{N}$. Indeed, both (9) and (10) would imply that $\Lambda_d(1) = \sup_{\ell^*_d\in\mathcal{L}_d} \rho_{\{\tilde{\ell}^*_d\}} = \sup_{\ell^*_d\in\mathcal{L}_d} \sup_{y\in\mathbb{N}} \rho_{\tilde{\ell}^*_d}(y) = \rho_{\tilde{\mathcal{L}}_d}$, that is, the claim.

We give a further lemma.

**Lemma 9.** We have $\Lambda_d(1) < \Lambda_{d+1}(1)$.

Proof of Lemma 9. As $\left\{ {d \atop j} \right\} \leq j\left\{ {d \atop j} \right\} + \left\{ {d \atop j-1} \right\} = \left\{ {d+1 \atop j} \right\}$ for any $j \in [d]$, we have

$$\begin{aligned}\Lambda_d(1) &= \sum_{j\in[d]}\left(\frac{j+2}{j+1}\right)\left\{ {d \atop j} \right\} \leq \sum_{j\in[d]}\left(\frac{j+2}{j+1}\right)\left\{ {d+1 \atop j} \right\} \\ &< \sum_{j\in[d+1]}\left(\frac{j+2}{j+1}\right)\left\{ {d+1 \atop j} \right\} = \Lambda_{d+1}(1),\end{aligned}$$

and the claim of the lemma follows. □

Let us fix an arbitrary $y \in \mathbb{N}$ and a latency function $\ell^*_d \in \mathcal{L}_d$, that is, $\ell^*_d(x) = \sum_{q=0}^{d} \alpha_q x^q$, for some coefficients $\alpha_0, \ldots, \alpha_d \geq 0$.

Let $\beta_q := \alpha_q \sum_{h\in[y]} h^q$ for any $q \in [d] \cup \{0\}$. We have

$$\Lambda_d(1) = \max_{q\in[d]\cup\{0\}} \Lambda_q(1) \geq \max_{q\in[d]\cup\{0\}} \Lambda_q(y)$$
$$= \max_{q\in[d]\cup\{0\}} \rho_{\tilde{\ell}_q}(y) \geq \frac{\sum_{q=0}^{d} \beta_q \cdot \rho_{\tilde{\ell}_q}(y)}{\sum_{q=0}^{d} \beta_q}$$
$$= \frac{\sum_{q=0}^{d} \alpha_q \left(\sum_{x=0}^{\infty} \sum_{h\in[x]} h^q \frac{y^x}{x!e^y}\right)}{\sum_{q=0}^{d} \left(\alpha_q \sum_{h\in[y]} h^q\right)}$$
$$= \frac{\sum_{x=0}^{\infty} \sum_{h\in[x]} \left(\sum_{q=0}^{d} \alpha_q h^q\right) \frac{y^x}{x!e^y}}{\sum_{h\in[y]} \left(\sum_{q=0}^{d} \alpha_q h^q\right)}$$
$$= \frac{\sum_{x=0}^{\infty} \left(\sum_{h\in[x]} \ell_d^*(h)\right) \frac{y^x}{x!e^y}}{\sum_{h\in[y]} \ell_d^*(h)} = \frac{\sum_{x=0}^{\infty} x\tilde{\ell}_d^*(x) \frac{y^x}{x!e^y}}{y\tilde{\ell}_d^*(y)} = \rho_{\tilde{\ell}_d^*}(y),$$

where the first and the second equality, respectively, follow from Lemma 9 and equality (8) (applied with $q$ in place of $d$). By the above inequalities, inequality (10) follows. Finally, because of the above observations, both (9) and (10) show the claim. □

**Remark 1** (A variant of Dobinski's formula)**.** We observe that Theorem 5 and, in particular, the equality shown in Lemma 7, is of independent interest, as it provides a variant of Dobinski's formula [24]. Indeed, Dobinski's formula states that the $d$-th Bell number $B_d = \sum_{j\in[d+1]} \left\{ {d \atop j} \right\}$ is equal to $\sum_{x=0}^{\infty} \frac{x^d}{x!e}$, while Lemma 7, applied with $y = 1$, states that $\sum_{j\in[d]} \left\{ {d \atop j} \right\} \left(\frac{j+2}{j+1}\right) = \sum_{x=0}^{\infty} \frac{\sum_{h\in[x]} h^d}{x!e}$.

## 5 MIN POTENTIAL WITH NON-INCREASING LATENCIES

This section is devoted to Min Potential for congestion games having *non-increasing* latency functions. We shall see that the situation significantly differs from the non-decreasing case. In particular, Proposition 1 does not hold for non-increasing latency functions.

**Proposition 3.** *If the game is symmetric and $S$ denotes the strategy space of every player,* Min Potential *can be solved in $|S|$ steps.*

Proof. Fix a symmetric congestion game $\mathcal{G}$ with non-increasing latency functions. Let us first observe that if two distinct strategies $a, b \in S$ are actually played in a state $\mathbf{s}$, then one of the following modifications of $\mathbf{s}$ gives a new state $\mathbf{s}'$ satisfying $\Phi_{\mathcal{G}}(\mathbf{s}') \leq \Phi_{\mathcal{G}}(\mathbf{s})$: either all the players playing $a$ change for $b$, or all the players playing $b$ change for $a$. (The other players stick to their strategy.)

Let $i$ (resp., $j$) be a player such that $s_i = a$ (resp., $s_j = b$). If $c_i(\mathbf{s}) \geq c_j(\mathbf{s})$, then all the players playing $a$ under $\mathbf{s}$ can change their strategy for $b$, and their individual cost will not increase. Indeed, the latency functions being non-increasing, the new cost of the deviating players would be at most $c_j(\mathbf{s})$. If $c_i(\mathbf{s}) < c_j(\mathbf{s})$, then all the players playing $b$ under $\mathbf{s}$ can change their strategy for $a$, and their individual cost will decrease since it will be at most $c_i(\mathbf{s})$ (the fact that every latency function is non-increasing is used again). Since Rosenthal's function is an exact potential, we deduce that $\Phi_{\mathcal{G}}(\mathbf{s}) \geq \Phi_{\mathcal{G}}(\mathbf{s}')$, where $\mathbf{s}'$ is the state obtained from $\mathbf{s}$ by grouping the players of $a$ and $b$ either onto $a$, or onto $b$.

We know from the above observation that there always exists a strategy profile $\mathbf{s}^*$ that minimizes $\Phi_{\mathcal{G}}(\mathbf{s}^*)$ in which all the players adopt the exact same strategy. From an algorithmic viewpoint, one can try every strategy $s \in S$, and retain the strategy profile $(s, \ldots, s)$ which minimizes Rosenthal's potential. □

Proposition 3 implies that Min Potential can be solved in $m^{O(1)}$ operations when every strategy consists of selecting a constant number of resources. By a reduction from vertex cover, the following result states that Min Potential is hard when the symmetry property is dropped, even if other parameters of the game are significantly restricted.

**Theorem 10.** Min Potential *is NP-hard, even if all the resources have the same latency function, and all the players only have two singleton strategies.*

To conclude, observe that the approximability of Min Potential when $\mathtt{size}(\mathcal{G}) = 1$ is similar to the approximability of set cover (every player is "covered" by her selected resource). Recall that $H_k := \sum_{i=1}^{k} \frac{1}{i} = \Theta(\ln k)$ denotes the $k$-th harmonic number.

**Theorem 11.** Min Potential *admits a $H_n$-approximation algorithm for singleton congestion games. Moreover, the approximation ratio $H_n$ is best possible unless P = NP.*

## 6 CONCLUSION

We have considered the complexity of building a state of minimum potential in congestion games with monotone latency functions. Our results show that the symmetry of the players' strategies, together with the maximum number of resources used simultaneously, plays an important role.

Although it is long known that, in general, computing a pure Nash equilibrium in a congestion game is PLS-complete [18], an intriguing question for future work is about the complexity of computing a pure Nash equilibrium (i.e., a *local* minimum of Rosenthal's potential instead of a global minimum) in a monotone non-decreasing congestion game with $\mathtt{size} = 2$ (general strategies) or $\mathtt{size} = 3$ (symmetric strategies). The same question is of interest in monotone non-increasing congestion games with $\mathtt{size} = 2$ (general strategies).

Natural dynamics like better or best response, starting from any initial state, always converge towards a pure Nash equilibrium in congestion games, and the time convergence is known to be polynomial if the instance is singleton [21], or the strategies are bases of a matroid [1]. An interesting question is to bound the worst-case convergence time of these dynamics for (possibly monotone) congestion games with $\mathtt{size} = O(1)$. We proposed an approximation algorithm for congestion games with non-increasing latency functions and $\mathtt{size} = 1$ (Theorem 11) but it would be interesting to have an approximation for bigger sizes.

## ACKNOWLEDGMENTS

Laurent Gourvès is supported by Agence Nationale de la Recherche (ANR), project THEMIS ANR-20-CE23-0018. Vittorio Bilò and Cosimo Vinci are supported by: the PNRR MIUR project FAIR - Future AI Research (PE00000013), Spoke 9 - Green-aware AI; MUR - PNRR IF Agro@intesa; the Project SERICS (PE00000014) under the NRRP MUR program funded by the EU – NGEU; GNCS-INdAM.

## REFERENCES


[1] Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. 2008. On the impact of combinatorial structure on congestion games. *Journal of the ACM* 55, 6 (2008), 25:1–25:22.

[2] Susanne Albers and Pascal Lenzner. 2013. On approximate Nash equilibria in network design. *Internet Mathematics* 9, 4 (2013), 384–405.

[3] Elliot Anshelevich, Anirban Dasgupta, Jon M. Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. 2008. The Price of Stability for Network Design with Fair Cost Allocation. *SIAM J. Comput.* 38, 4 (2008), 1602–1623.

[4] Ronald D. Armstrong and Zhiying Jin. 1997. A new strongly polynomial dual network simplex algorithm. *Mathematical Programming* 78 (1997), 131–148.

[5] Daniel Berend and Tamir Tassa. 2000. Improved Bounds on Bell Numbers and on Moments of Sums of Random Variables. *Probability and Mathematical Statistics* 30, 2 (2000), 185–205.

[6] Vittorio Bilò, Angelo Fanelli, Laurent Gourvès, Christos Tsoufis, and Cosimo Vinci. 2024. Minimizing Rosenthal's Potential in Monotone Congestion Games. arXiv:2408.11489 [cs.GT] https://arxiv.org/abs/2408.11489

[7] Vittorio Bilò and Cosimo Vinci. 2019. Dynamic Taxes for Polynomial Congestion Games. *ACM Transantions on Economics and Computation* 7, 3 (2019), 15:1–15:36.

[8] Vittorio Bilò and Cosimo Vinci. 2023. *Coping with Selfishness in Congestion Games: Analysis and Design via LP Duality.* Springer, Berlin, Heidelberg.

[9] Vittorio Bilò and Cosimo Vinci. 2024. Enhancing the Efficiency of Altruism and Taxes in Affine Congestion Games through Signalling. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 9 (2024), 9511–9518.

[10] Vittorio Bilò, Michele Flammini, Gianpiero Monaco, and Luca Moscardelli. 2021. Computing approximate Nash equilibria in network congestion games with polynomially decreasing cost functions. *Distributed Computing* 34, 1 (2021), 1–14.

[11] Ioannis Caragiannis, Michele Flammini, Christos Kaklamanis, Panagiotis Kanellopoulos, and Luca Moscardelli. 2011. Tight bounds for selfish and greedy load balancing. *Algorithmica* 61, 3 (2011), 606–637.

[12] Ioannis Caragiannis, Christos Kaklamanis, and Panagiotis Kanellopoulos. 2010. Taxes for Linear Atomic Congestion Games. *ACM Transactions on Algorithms* 7, 1, Article 13 (2010), 31 pages.

[13] Steve Chien and Alistair Sinclair. 2011. Convergence to approximate Nash equilibria in congestion games. *Games and Economic Behavior* 71, 2 (2011), 315–327.

[14] George Christodoulou and Elias Koutsoupias. 2005. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Proceedings of the 13th Annual European Symposium on Algorithms (ESA)*. Springer Berlin Heidelberg, Berlin, Heidelberg, 59–70.

[15] George Christodoulou, Vahab S. Mirrokni, and Anastasios Sidiropoulos. 2012. Convergence and approximation in potential games. *Theoretical Computer Science* 438 (2012), 13–27.

[16] Alberto Del Pia, Michael Ferris, and Carla Michini. 2017. Totally Unimodular Congestion Games. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Society for Industrial and Applied Mathematics, USA, 577–588.

[17] Eyal Even-Dar, Alexander Kesselman, and Yishay Mansour. 2007. Convergence time to Nash equilibrium in load balancing. *ACM Transactions on Algorithms* 3, 3 (2007), 1–32.

[18] Alex Fabrikant, Christos H. Papadimitriou, and Kunal Talwar. 2004. The complexity of pure Nash equilibria. In *Proceedings of the 36th ACM Symposium on Theory of Computing (STOC)*. Association for Computing Machinery, New York, NY, USA, 604–612.

[19] Dimitris Fotakis. 2010. Congestion Games with Linearly Independent Paths: Convergence Time and Price of Anarchy. *Theory of Computing Systems* 47, 1 (2010), 113–136.

[20] Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. 1989. *Concrete mathematics - a foundation for computer science.* Addison-Wesley, Reading, MA.

[21] Samuel Ieong, Robert McGrew, Eugene Nudelman, Yoav Shoham, and Qixiang Sun. 2005. Fast and Compact: A Simple Class of Congestion Games. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI)*. AAAI Press, Pittsburgh, Pennsylvania, 489–494.

[22] Pieter Kleer and Guido Schäfer. 2021. Computation and efficiency of potential function minimizers of combinatorial congestion games. *Mathematical Programming* 190, 1 (2021), 523–560.

[23] Elias Koutsoupias and Christos H. Papadimitriou. 1999. Worst-case Equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*. Springer Berlin Heidelberg, Berlin, Heidelberg, 404–413.

[24] Toufik Mansour and Matthias Schork. 2015. *Commutation Relations, Normal Ordering, and Stirling Numbers.* Chapman and Hall/CRC, Boca Raton, FL.

[25] Carol A. Meyers and Andreas S. Schulz. 2012. The complexity of welfare maximization in congestion games. *Networks* 59, 2 (2012), 252–260.

[26] Dov Monderer and Lloyd S. Shapley. 1996. Potential Games. *Games and Economic Behavior* 14 (1996), 124–143.

[27] Dario Paccagnan, Rahul Chandan, Bryce L. Ferguson, and Jason R. Marden. 2021. Optimal Taxes in Atomic Congestion Games. *ACM Transactions on Economics and Computation* 9, 3 (2021), 19:1–19:33.

[28] Dario Paccagnan and Martin Gairing. 2024. In Congestion Games, Taxes Achieve Optimal Approximation. *Operations Research* 72, 3 (2024), 966–982.

[29] B. C. Rennie and A. J. Dobson. 1969. On Stirling numbers of the second kind. *Journal of Combinatorial Theory* 7, 2 (1969), 116–121.

[30] Robert W. Rosenthal. 1973. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory* 2 (1973), 65–67.

[31] Walter Rudin. 1976.. *Principles of mathematical analysis* (3rd ed.). McGraw-Hill,, New York, NY.

[32] Alexander Schrijver. 2003. *Combinatorial Optimization: Polyhedra and Efficiency.* Number v. 1 in Algorithms and Combinatorics. Springer-Verlag, Berlin, Heidelberg. https://books.google.fr/books?id=mqGeSQ6dJycC

[33] Vasilis Syrgkanis. 2010. The complexity of equilibria in cost sharing games. In *Proceedings of the 6th Workshop on Internet and Network Economics (WINE)*. Springer Berlin Heidelberg, Berlin, Heidelberg, 366–377.

[34] Vipin Ravindran Vijayalakshmi and Alexander Skopalik. 2020. Improving Approximate Pure Nash Equilibria in Congestion Games. In *Proceedings of the 16th International Conference on Web and Internet Economics (WINE)*. Springer-Verlag, Berlin, Heidelberg, 280–294.

[35] Cosimo Vinci. 2019. Non-atomic one-round walks in congestion games. *Theoretical Computer Science* 764 (2019), 61–79. Selected papers of ICTCS 2016 (The Italian Conference on Theoretical Computer Science (ICTCS).

[36] Berthold Vöcking. 2006. Congestion Games: Optimization in Competition. In *Algorithms and Complexity in Durham 2006 - Proceedings of the Second ACiD Workshop (Texts in Algorithmics, Vol. 7)*, Hajo Broersma, Stefan S. Dantchev, Matthew Johnson, and Stefan Szeider (Eds.). King's College, London, UK, 9–20.

[37] Scott Webster and Kenneth R. Baker. 1995. Scheduling Groups of Jobs on a Single Machine. *Operations Research* 43, 4 (1995), 692–703. https://doi.org/10.1287/opre.43.4.692