

Nash Social Welfare in Selfish and Online Load Balancing (Short Paper)*

Vittorio Bilò^{1,†}, Gianpiero Monaco^{2,†}, Luca Moscardelli^{3,†} and Cosimo Vinci^{4,*,†}

¹Università del Salento, Italy

²Università degli Studi dell'Aquila, Italy

³Università degli Studi "Gabriele D'Annunzio" Chieti - Pescara, Italy

⁴Università degli Studi di Salerno, Italy

Abstract

In load balancing problems there is a set of clients, each wishing to select a resource from a set of permissible ones, in order to execute a certain task. Each resource has a latency function, which depends on its workload, and a client's cost is the completion time of her chosen resource. Two fundamental variants of load balancing problems are *selfish load balancing* (aka. *load balancing games*), where clients are non-cooperative selfish players aimed at minimizing their own cost solely, and *online load balancing*, where clients appear online and have to be irrevocably assigned to a resource without any knowledge about future requests. We revisit both problems under the objective of minimizing the *Nash Social Welfare*, i.e., the geometric mean of the clients' costs. To the best of our knowledge, despite being a celebrated welfare estimator in many social contexts, the Nash Social Welfare has not been considered so far as a benchmarking quality measure in load balancing problems. We provide tight bounds on the price of anarchy of pure Nash equilibria and on the competitive ratio of the greedy algorithm under very general latency functions, including polynomial ones. For this particular class, we also prove that the greedy strategy is optimal, as it matches the performance of any possible online algorithm.

Keywords

Load Balancing, Nash Welfare, Nash Equilibria, Price of Anarchy, Online Algorithms

1. Introduction

In load balancing problems there is a set of clients, each wishing to select a resource from a set of permissible ones, in order to execute a certain task. Each resource has a latency function, which depends on its workload, and a client's cost is the completion time of her chosen resource. These problems stand at the foundations of the Theory of Computing and have been studied

Proceedings of the 23rd Italian Conference on Theoretical Computer Science, Rome, Italy, September 7-9, 2022

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

Extended versions of this paper appear in the proceedings of "WINE 2020" [1] and in "ACM Transactions on Economics and Computation" [2].


*Corresponding author.

† These authors contributed equally.

✉ vittorio.bilo@unisalento.it (V. Bilò); gianpiero.monaco@univaq.it (G. Monaco); luca.moscardelli@unich.it (L. Moscardelli); cvinci@unisa.it (C. Vinci)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

under a variety of objective functions, such as the maximum client's cost (aka. the makespan) [3, 4, 5, 6] and the average weighted client's cost (see [7] for an excellent survey).

Two extensively studied variants of load balancing problems are *selfish load balancing* [8] (aka. *load balancing games*) and *online load balancing* [3]. Selfish load balancing, where clients are non-cooperative selfish players aimed at minimizing their own cost solely, constitutes a notable subclass of *weighted congestion games* [9] and, as such, enjoys some nice theoretical properties. For instance, they always admit pure Nash Equilibria [10]. In online load balancing, instead, clients appear online and have to be irrevocably assigned to a resource without any knowledge about future requests. Interpreting the set of clients of a load balancing problem as a society and adopting the terminology of welfare economics, the makespan and the average weighted client's cost objective functions get called, respectively, the *Egalitarian Social Welfare* (ESW) and the *Utilitarian Social Welfare* (USW). In the case of unweighted tasks, the ESW is defined as $\max_i x_i$, and the USW is defined as $\frac{1}{n} \sum_i x_i$, where n is the number of clients and $x = (x_1, x_2, \dots)$ is the vector encoding the clients' costs. Another interesting social function is the *Nash Social Welfare* (NSW) [11], which is defined as $(\prod_i x_i)^{\frac{1}{n}}$, i.e., as the geometric mean of the clients' costs. These definitions naturally extend to the more general case of weighted tasks.

The NSW is a celebrated welfare measure in many settings, such as Fisher markets [12, 13] and fair division [14, 15], as it satisfies a set of interesting properties and achieves a balanced compromise between the equity of the egalitarian social welfare function and the efficiency of the utilitarian one. Despite being an appealing welfare estimator in many social contexts, the Nash Social Welfare has not been considered so far as a benchmarking quality measure in load balancing or general cost minimization problems. Anyway, the preference relation among different outcomes induced by the NSW (where an outcome is better if its NSW is lower) satisfies several interesting properties.

An interesting motivation for considering the NSW in load balancing comes from the following observation. An alternative reasonable way to define a client's cost can come by taking the ratio between the completion time of her chosen resource and the completion time she could obtain when being the only client in the system (i.e., when she is the unique user of the fastest resource). This definition avoids situations where the cost of a specific client determines almost completely the value of the social welfare. This happens, for instance, when there is a client i owing a highly time-consuming task. Here, both the utilitarian and the egalitarian social welfare end up depending on the cost of i , thus almost neglecting the other clients' costs. In this setting, the NSW is the proper metric to use. More generally, the NSW is the only correct mean to use when averaging normalized results, that is, results that are presented as ratios to reference values [16]. In light of the above example, it is important to emphasize the property of *player-specific scale-independence* of the NSW in load balancing, stating that the preference relation between two outcomes does not change if the costs of different players are scaled by player-specific values in both outcomes. We point out that such property, in general, does not hold for the ESW and the USW, while it is always satisfied by the NSW.

The preference relation induced by the NSW in load balancing also satisfies the *Pareto optimality*, prescribing that an outcome in which the cost of a player improves while all other costs do not get worse has to be preferred.

To conclude, the effectiveness of the Nash Social Welfare as quality measure does not hold for

profit maximization settings only, but also for load balancing and more general cost minimization frameworks, where the lower the NSW, the better the considered outcome.

2. Our Contribution

In this work, we revisit both selfish and online load balancing under the objective of minimizing the NSW. To the best of our knowledge, this is the first work adopting the NSW as a benchmarking quality measure in load balancing problems. Indeed, the performances of Nash equilibria in load balancing games, as well as the competitive ratio for online load balancing, have been widely studied under the USW and the ESW, but never under the NSW. Furthermore, most of the literature on NSW is about the problem of allocating a set of items among players with the aim of maximizing the NSW [17, 18, 14, 19, 15, 20], while in this work the NSW is considered as a quality measure to be minimized.

We analyze the price of anarchy [21] of pure Nash equilibria (the loss in optimality due to selfish behavior) and the competitive ratio of online algorithms (the loss in optimality due to lack of information) under very general latency functions. These questions have been widely addressed under the USW¹ and the ESW², but never under the NSW.

We notice that, by adopting the NSW as a new metric, we are not going to modify the set of Nash equilibria, but only their social values. The main difference between the NSW and the classical notion of USW consists in the fact that, while in the latter the players' costs are summed, in the former they are multiplied. This may lead to think that, by turning the costs into their logarithms, a classical utilitarian analysis can be easily adapted to deal with the NSW. Actually, this is not the case. In fact, on the one hand, using this idea for bounding a performance ratio (e.g., the price of anarchy or the competitive ratio), one obtains a bound on the ratio between two logarithms (each one having the product of the players' costs as argument). On the other hand, we are interested in bounding the ratio between the argument of these logarithms, and there is no direct correlation between these two ratios (notice that logarithm of the latter ratio is equal to the difference between the corresponding utilitarian social costs, and therefore it is not related to the former one). Thus, the analysis of the NSW requires different proof arguments. In order to have another evidence of this fact, it is worth noticing that the results obtained for the NSW substantially differ from the ones holding for the USW, not only from a quantitative point of view, but also from a qualitative one. In fact, while it is well known (see [39]) that for the USW the simpler combinatorial structure of load balancing games does not improve the price of anarchy of general congestion games, we show that, for the NSW the price of anarchy drops from n to 2, even for the case of linear latency functions.

All upper bounds shown in this paper are quite general, given that they hold for any family of non-decreasing and positive latency functions. Moreover, the provided matching lower bounds hold for latency functions verifying mild assumptions; it is worth to remark that they are satisfied by the well studied class of polynomial latency functions, and by many other ones.

¹Relatively to the USW, the interested reader can refer to [22, 23, 24, 25, 26, 27, 28, 29] for the price of anarchy, and to [30, 31, 24, 32, 33, 34, 35] for the competitive ratio of online algorithms.

²Relatively to the ESW, the interested reader can refer to [21, 36] for the price of anarchy, and to [37, 38] for the competitive ratio of online algorithms.

In particular, the following theorem provides an upper bound to the price of anarchy for the case of weighted load balancing games:

Theorem 1. *Let \mathcal{C} be a class of latency functions. The price of anarchy $\text{PoA}_W(\mathcal{C})$ (w.r.t. the NSW) of weighted load balancing games with latencies in \mathcal{C} is*

$$\text{PoA}_W(\mathcal{C}) \leq \sup_{\substack{f_1, f_2 \in \mathcal{C}, \\ k_1, k_2, o_1, o_2 \in \mathbb{R}: \\ k_1 \geq o_1 > 0, o_2 > k_2 \geq 0}} \left(\frac{f_1(k_1 + o_1)}{f_1(o_1)} \right)^{\frac{(o_2 - k_2)o_1}{k_1 o_2 - k_2 o_1}} \left(\frac{f_2(k_2 + o_2)}{f_2(o_2)} \right)^{\frac{(k_1 - o_1)o_2}{k_1 o_2 - k_2 o_1}}.$$

Furthermore, we show that the above upper bound is tight under mild assumptions.

Similarly, we focus on unweighted games (a special case of weighted ones) by providing tight bounds that, in general, are lower than the ones that can be obtained for weighted games. In the following theorem we provide an upper bound for unweighted games, that is tight under mild assumptions.

Theorem 2. *Let \mathcal{C} be a class of latency functions. The price of anarchy $\text{PoA}_U(\mathcal{C})$ (w.r.t. the NSW) of unweighted load balancing games with latencies in \mathcal{C} is*

$$\text{PoA}_U(\mathcal{C}) \leq \sup_{f \in \mathcal{C}, k \in \mathbb{N}, o \in [k]} \left(\frac{f(k+1)}{f(o)} \right)^{\frac{o}{k}}.$$

We also show that, when considering polynomial latency functions of degree p , the analyses for weighted and unweighted games give the same tight bound of 2^p . Furthermore, when considering weighted games, the tight bound of 2^p holds even for symmetric games and for games with identical resources.

We also provide a tight analysis holding for non-atomic games, as stated in the following theorem:

Theorem 3. *Let \mathcal{C} be a class of latency functions. The price of anarchy $\text{PoA}_N(\mathcal{C})$ (w.r.t. the NSW) of non-atomic load balancing games with latencies in \mathcal{C} is*

$$\text{PoA}_N(\mathcal{C}) \leq \sup_{f \in \mathcal{C}, k, o \in \mathbb{R}: k \geq o > 0} \left(\frac{f(k)}{f(o)} \right)^{\frac{o}{k}}.$$

A tight lower-bound, under mild assumptions, is attained by a simple Pigou-like network [40] (as well as for the utilitarian social welfare [27]); for the case of polynomial latency functions of degree p , we show that the price of anarchy is $\left(e^{\frac{1}{e}}\right)^p \simeq (1.44)^p$.

For the online setting, we analyze the greedy algorithm that assigns every client to a resource minimizing the total cost of the instance revealed up to the time of its appearance. We provide a tight analysis of the competitive ratio of the greedy algorithm, and we further show that, when considering polynomial latency functions of degree p , there exists no online algorithm achieving a competitive ratio better than the one of the greedy algorithm, that is equal to 4^p .

In Table 1, we consider the case of polynomial latency functions for both the price of anarchy and the competitive ratio, and we compare the performance under the NSW with that under the USW studied in some previous works.

	NSW	USW
Weighted (PoA)	2^p	$(\Phi_p)^{p+1} \sim \Theta\left(\frac{p}{\log(p)}\right)^{p+1}$, [22]
Unweighted (PoA)	2^p	$\frac{(k+1)^{2p+1} - k^{p+1}(k+2)^p}{(k+1)^{p+1} - (k+2)^p + (k+1)^p - k^{p+1}} \sim \Theta\left(\frac{p}{\log(p)}\right)^{p+1}$, [22]
Non-atomic (PoA)	$\left(e^{\frac{1}{e}}\right)^p$	$(1 - p(p+1)^{-(p+1)/p})^{-1} \sim \Theta\left(\frac{p}{\log(p)}\right)$, [27]
Online (CR)	4^p	$(2^{1/(p+1)} - 1)^{-(p+1)} \sim \Theta(p)^{p+1}$, [32]

Table 1

Tight bounds on the performance of load balancing with polynomial latency functions of maximum degree p , under the NSW and the USW. Φ_p denotes the unique solution of equation $x^{p+1} = (x+1)^p$, and $k := \lfloor \Phi_p \rfloor$. We observe that the performance under the NSW case is definitely better (even asymptotically) than that under the USW case, except for the non-atomic setting.

3. Future Works

Our work leaves several research directions. Relatively to the the optimization problem of minimizing the NSW in weighted load balancing, it would be good to show better approximation factors than that provided by the greedy algorithm (whose optimality is guaranteed in the online setting only). A further research direction is considering the NSW for other classes or variants of load balancing and congestion games (e.g., [41, 42, 43, 44, 45, 46, 47]), and analyzing the resulting efficiency. It would be also important to consider other quality metrics to evaluate the performance under the NSW (e.g., the price of stability [41, 48, 49, 50, 51, 52]). Finally, in light of the suboptimality of the NSW in selfish load balancing, it would be nice to design and analyze mechanisms to improve the social performance (e.g., taxing mechanisms [53, 54], Stackelberg strategies [55, 25, 56], coordination mechanisms [57, 58, 59]).

References

- [1] V. Bilò, G. Monaco, L. Moscardelli, C. Vinci, Nash social welfare in selfish and online load balancing, in: Proceedings of the 16th International Conference on Internet and Network Economics (WINE), 2020, pp. 323–337.
- [2] V. Bilò, G. Monaco, L. Moscardelli, C. Vinci, Nash social welfare in selfish and online load balancing, ACM Trans. Econ. Comput. (2022). Accepted.
- [3] R. L. Graham, Bounds for certain multiprocessing anomalies, The Bell System Technical Journal 45 (1966) 1563–1581.
- [4] D. S. Hochbaum, D. B. Shmoys, Using dual approximation algorithms for scheduling problems: theoretical and practical results, Journal of ACM 34 (1987) 144–162.
- [5] E. Horowitz, S. K. Sahni, Exact and approximate algorithms for scheduling nonidentical processors, Journal of ACM 23 (1976) 317–327.
- [6] J. K. Lenstra, D. B. Shmoys, E. Tardos, Approximation algorithms for scheduling unrelated parallel machines, Mathematical Programming 46 (1990) 259–271.
- [7] C. Chekuri, S. Khanna, Handbook of Scheduling: Algorithms, Models, and Performance Analysis, Chapman & Hall/CRC, 2004.

- [8] B. Vöcking, *Algorithmic Game Theory*, Cambridge, 2007, pp. 517–542.
- [9] R. W. Rosenthal, A class of games possessing pure-strategy Nash equilibria, *International Journal of Game Theory* 2 (1973) 65–67.
- [10] S. Ieong, R. McGrew, E. Nudelman, Y. Shoham, Q. Sun, Fast and compact: A simple class of congestion games, in: *Proceedings of the 20th AAAI Conference on Artificial Intelligence (AAAI)*, 2005, pp. 489–494.
- [11] J. Nash, The bargaining problem, *Econometrica* 18 (1950) 155–162.
- [12] X. Bei, J. Garg, M. Hoefer, K. Mehlhorn, Earning and utility limits in fisher markets, *ACM Transactions on Economics and Computation* 7 (2019) 10:1–10:35.
- [13] W. C. Brainard, H. E. Scarf, How to compute equilibrium prices in 1891, *Cowles Foundation Discussion Paper* 1270 (2000).
- [14] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, J. Wang, The unreasonable fairness of maximum Nash welfare, in: *Proceedings of the 2016 ACM Conference on Economics and Computation (EC)*, 2016, pp. 305–322.
- [15] R. Cole, V. Gkatzelis, Approximating the Nash social welfare with indivisible items, *SIAM J. Comput.* 47 (2018) 1211–1236.
- [16] P. J. Fleming, J. Wallace, How not to lie with statistics: The correct way to summarize benchmark results, *Commun. ACM* 29 (1986) 218–221.
- [17] S. Brânzei, V. Gkatzelis, R. Mehta, Nash social welfare approximation for strategic agents, in: *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*, 2017, pp. 611–628.
- [18] I. Caragiannis, N. Gravin, X. Huang, Envy-freeness up to any item with high Nash welfare: The virtue of donating items, in: *Proceedings of the 2019 ACM Conference on Economics and Computation (EC)*, 2019, pp. 527–545.
- [19] R. Cole, N. R. Devanur, V. Gkatzelis, K. Jain, T. Mai, V. V. Vazirani, S. Yazdanbod, Convex program duality, fisher markets, and nash social welfare, in: *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*, 2017, pp. 459–460.
- [20] J. Garg, M. Hoefer, K. Mehlhorn, Approximating the Nash social welfare with budget-additive valuations, in: *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2018, pp. 2326–2340.
- [21] E. Koutsoupias, C. Papadimitriou, Worst-case equilibria, in: *Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science (STACS)*, 1999, pp. 404–413.
- [22] S. Aland, D. Dumrauf, M. Gairing, B. Monien, F. Schoppmann, Exact price of anarchy for polynomial congestion games, *SIAM Journal on Computing* 40 (2011) 1211–1233.
- [23] K. Bhawalkar, M. Gairing, T. Roughgarden, Weighted congestion games: price of anarchy, universal worst-case examples, and tightness, *ACM Transactions on Economics and Computation* 2 (2014) 1–23.
- [24] V. Bilò, C. Vinci, On the impact of singleton strategies in congestion games, in: *Proceedings of the 25th Annual European Symposium on Algorithms, (ESA)*, 2017, pp. 17:1–17:14.
- [25] D. Fotakis, Stackelberg strategies for atomic congestion games, *Theor. Comp. Sys.* 47 (2010) 218–249.
- [26] M. Gairing, F. Schoppmann, Total latency in singleton congestion games, in: *Proceedings of the Third International Workshop on Internet and Network Economics (WINE)*, volume 4858 of *LNCS*, 2007, pp. 381–387.

- [27] T. Roughgarden, The price of anarchy is independent of the network topology, *J. Comput. Syst. Sci.* 67 (2003) 341–364.
- [28] T. Roughgarden, E. Tardos, How bad is selfish routing?, *J. ACM* 49 (2002) 236–259.
- [29] T. Roughgarden, E. Tardos, Bounding the inefficiency of equilibria in nonatomic congestion games, *Games and Economic Behavior* 47 (2004) 389–403.
- [30] B. Awerbuch, A. Yossi, E. F. Grove, M. Kao, P. Krishnan, J. S. Vitter, Load balancing in the l_p norm, in: *Proceedings of the 36th Annual Symposium on Foundations of Computer Science (FOCS)*, 1995, pp. 383–391.
- [31] V. Bilò, A. Fanelli, M. Flammini, L. Moscardelli, Performances of one-round walks in linear congestion games, *Theory of Computing Systems* 49 (2011) 24–45.
- [32] I. Caragiannis, Better bounds for online load balancing on unrelated machines, in: *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2008, pp. 972–981.
- [33] G. Christodoulou, V. S. Mirrokni, A. Sidiropoulos, Convergence and approximation in potential games, *Theor. Comput. Sci.* 438 (2012) 13–27.
- [34] M. Klimm, D. Schmand, A. Tönnis, The online best reply algorithm for resource allocation problems, in: *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT)*, 2019, pp. 200–215.
- [35] C. Vinci, Non-atomic one-round walks in congestion games, *Theor. Comput. Sci.* 764 (2019) 61–79.
- [36] A. Czumaj, B. Vöcking, Tight bounds for worst-case equilibria, *ACM Trans. Algorithms* 3 (2007) 4:1–4:17.
- [37] J. Aspnes, Y. Azar, A. Fiat, S. A. Plotkin, O. Waarts, On-line routing of virtual circuits with applications to load balancing and machine scheduling, *Journal of ACM* 44 (1997) 486–504.
- [38] Y. Azar, J. Naor, R. Rom, The competitiveness of on-line assignments, in: *Proceedings of the Third Annual ACM/SIGACT-SIAM Symposium on Discrete Algorithms (SODA)*, 1992, pp. 203–210.
- [39] I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, L. Moscardelli, Tight bounds for selfish and greedy load balancing, *Algorithmica* 61 (2011) 606–637.
- [40] A. C. Pigou, *The economics of welfare*, London: Macmillan and Co., 1938.
- [41] E. Anshelevich, A. Dasgupta, J. M. Kleinberg, É. Tardos, T. Wexler, T. Roughgarden, The price of stability for network design with fair cost allocation, *SIAM J. Comput.* 38 (2008) 1602–1623.
- [42] F. Benita, V. Bilò, B. Monnot, G. Piliouras, C. Vinci, Data-driven models of selfish routing: Why price of anarchy does depend on network topology, in: *Proceedings of the 16th International Conference on Internet and Network Economics (WINE)*, 2020, pp. 252–265.
- [43] V. Bilò, M. Flammini, V. Gallotti, C. Vinci, On multidimensional congestion games, *Algorithms* 13 (2020) 261.
- [44] V. Bilò, L. Moscardelli, C. Vinci, Uniform mixed equilibria in network congestion games with link failures, in: *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP)*, 2018, pp. 146:1–146:14.
- [45] V. Bilò, C. Vinci, The price of anarchy of affine congestion games with similar strategies, *Theor. Comput. Sci.* 806 (2020) 641–654.
- [46] J. Correa, J. de Jong, B. de Keijzer, M. Uetz, The curse of sequentiality in routing games, in:

- Proceedings of the 11th International Conference on Web and Internet Economics (WINE), volume 9470 of *LNCS*, 2015, pp. 258–271.
- [47] D. Fotakis, V. Gkatzelis, A. C. Kaporis, P. G. Spirakis, The impact of social ignorance on weighted congestion games, in: Proceedings of the 5th International Conference on Internet and Network Economics (WINE), 2009, pp. 316–327.
 - [48] V. Bilò, I. Caragiannis, A. Fanelli, G. Monaco, Improved lower bounds on the price of stability of undirected network design games, *Theory Comput. Syst.* 52 (2013) 668–686.
 - [49] G. Christodoulou, M. Gairing, Price of stability in polynomial congestion games, *ACM Trans. Economics and Comput.* 4 (2016) 10:1–10:17.
 - [50] G. Christodoulou, M. Gairing, Y. Giannakopoulos, P. G. Spirakis, The price of stability of weighted congestion games, *SIAM J. Comput.* 48 (2019) 1544–1582.
 - [51] A. Fanelli, D. Leniowski, G. Monaco, P. Sankowski, The ring design game with fair cost allocation, *Theor. Comput. Sci.* 562 (2015) 90–100.
 - [52] V. Bilò, M. Flammini, L. Moscardelli, The price of stability for undirected broadcast network design with fair cost allocation is constant, *Games Econ. Behav.* 123 (2020) 359–376.
 - [53] V. Bilò, C. Vinci, Dynamic taxes for polynomial congestion games, *ACM Trans. Economics and Comput.* 7 (2019) 15:1–15:36.
 - [54] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, Taxes for linear atomic congestion games, *ACM Trans. Algorithms* 7 (2010) 13:1–13:31.
 - [55] V. Bilò, C. Vinci, On stackelberg strategies in affine congestion games, *Theory of Computing Systems* 63 (2019) 1228–1249.
 - [56] T. Roughgarden, Stackelberg scheduling strategies, *SIAM J. Comput.* 33 (2004) 332–350.
 - [57] I. Caragiannis, V. Gkatzelis, C. Vinci, Coordination mechanisms, cost-sharing, and approximation algorithms for scheduling, in: *Web and Internet Economics - 13th International Conference, (WINE)*, 2017, pp. 74–87.
 - [58] G. Christodoulou, K. Mehlhorn, E. Pyrga, Improving the price of anarchy for selfish routing via coordination mechanisms, in: *Proceedings of the 19th European Conference on Algorithms, ESA*, 2011, pp. 119–130.
 - [59] R. Cole, J. Correa, V. Gkatzelis, V. Mirrokni, N. Olver, Decentralized utilitarian mechanisms for scheduling games, *Games and Economic Behavior* 92 (2014) 306–326.