



HAL
open science

Selfish Transportation Games

Dimitris Fotakis, Laurent Gourvès, Jérôme Monnot

► **To cite this version:**

Dimitris Fotakis, Laurent Gourvès, Jérôme Monnot. Selfish Transportation Games. 43rd International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM 2017), Jan 2017, Limerick, Ireland. pp.176-187, 10.1007/978-3-319-51963-0_14 . hal-02104866

HAL Id: hal-02104866

<https://hal.science/hal-02104866v1>

Submitted on 19 Apr 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Selfish Transportation Games

Dimitris Fotakis¹, Laurent Gourvès², and Jérôme Monnot²

¹ National Technical University of Athens, 15780 Athens, Greece

² Université Paris-Dauphine, PSL Research University, CNRS UMR [7243], LAMSADE, 75016 Paris, France
fotakis@cs.ntua.gr, laurent.gourves@dauphine.fr, jerome.monnot@dauphine.fr

Abstract. We study a natural strategic situation arising from the selection of shared means of transportation. Some clients (the players) are located on different nodes of a given graph and they want to be transported from their location to a common destination point (*e.g.* school, airport). A fixed number of resources (also called buses) is available and each client has to choose exactly one. Individual costs depend on the route chosen by the buses and the distance between the nodes. We investigate the case where each bus has a static permutation which prescribes the order by which the clients are visited. We identify the cases admitting a pure strategy equilibrium and consider the construction of an equilibrium, via a dedicated algorithm, or a dynamics. We also determine the price of anarchy and the price of stability for two natural social functions.

Keywords: resource allocation game, existence and computation of equilibria, price of anarchy/stability

1 Introduction

In many applications some entities compete for the use of shared resources (*e.g.* processors, storage). These resources are typically rare enough to prevent the existence of an ideal situation where every entity is fully satisfied with the resources that it holds. Allocating scarce resources to a pool of agents is an important problem in the AGT community. An allocation can be found by a central planner who strives to optimize a prescribed social choice function. In practice, this approach is often too rigid as some users, who disregard the social cost and focus on their own individual cost, may not trust the planner's global solution. Therefore, one may prefer a more flexible mechanism, *i.e.* a decentralized way for constructing an allocation with which the agents can interact and be proactive. Unfortunately such a mechanism would fail if the agents were unable to commit in a feasible allocation of the resources, or if every allocation that the agents build has a high social cost. As a consequence, it is necessary to evaluate the worth of a mechanism with respect to at least two criteria: (*Stability*) Can the agents reach a stable state where no one is capable to acquire more revenue? (*Performance*) Are these stable states good from a social viewpoint?

We address these questions in a context where some agents, located on different places of a map, compete for the use of some public means of transportation (*e.g.* buses), so as to reach a common destination (*e.g.* an airport). Our work, motivated by services like *Dial-a-Ride*, *Lyft* or *Uber*, uses strategic games to model the situation. The models are called *selfish transportation games* because every agent (also called player) controls a part of the entire solution (*i.e.* which transportation means she decides to take) and this agent's choice is solely guided by her individual cost. Our study focuses on the existence of pure strategy equilibria, together with a worst-case analysis of the performance of the best (resp., worst) pure Nash equilibrium compared to configurations with minimum social cost. Two natural, yet different, notions of social cost are used. More importantly we concentrate on mechanisms where each bus visits and picks up the clients according to a fixed order (possibly different for each bus).

1.1 The Model

Let $G = (V, E)$ be an undirected graph with a source s and a destination t . The graph is also endowed with a distance function $d : V \times V \rightarrow \mathbb{R}_+$ which is possibly *metric*, *i.e.*, symmetric and it obeys the

triangle inequality. The *transportation game* has a set N of n players, and each one is located on a vertex of V . The goal of each player is to be transported from its location to the destination t at the lowest cost. There is a set M of $m \geq 2$ resources (also called buses). Each bus follows a path that starts from s , visits some players at their location, and finally reaches the destination t . We suppose that each bus j has its own algorithm \mathcal{A}_j which, given $V' \subseteq V$, determines its route, *i.e.* an $s - t$ path whose set of intermediate nodes is V' . Every algorithm \mathcal{A}_j is public. It is assumed that a bus always takes the direct link with distance $d(a, b)$ between two consecutive clients a and b and all links between two distinct clients are possible.

We consider a strategic game in which each player chooses by which bus it is picked. Thus, M is the strategy space of every player in N . There are different ways to define $c_i(\sigma)$, the individual cost of a player i under strategy profile σ . In this work $c_i(\sigma)$ is the distance travelled by σ_i (the bus selected by player i) between the original location of i (when player i is picked) and the destination t .

We suppose that each \mathcal{A}_j , for $j \in M$, is based on a permutation $\pi_j : \{1, \dots, n\} \rightarrow N$ (independent of the current strategy profile). Actually π_j indicates the reverse order by which the players are picked. This picking order is never violated, even if a bus visits a player's location more than once. The permutation is an expression of preferences, or priorities, that a resource has over the set of players (or their locations).

Example 1. Consider an instance with 4 players and 2 buses. The permutations of the buses are (1 2 3 4) and (1 4 2 3), respectively. Suppose player 1 chooses bus 1 whereas the others choose bus 2. Thus, bus 1 starts from s , visits 1 and goes to t . Bus 2 starts from s , visits players 3, 2, 4 and goes to t . Individual costs are $c_1(\sigma) = d(1, t)$, $c_4(\sigma) = d(4, t)$, $c_2(\sigma) = d(2, 4) + c_4(\sigma)$ and $c_3(\sigma) = d(3, 2) + c_2(\sigma)$.

In a totally equivalent model, every player is at t and wants to reach her location (*e.g.* airport to home). Each permutation π_j indicates by which order the players are dropped. In Example 1, bus 1 transports player 1 to his home. Bus 2 starts from t , drops player 4, then player 2, then player 3.

1.2 Motivation and Related Work

Transportation problems have a prominent place in operations research and combinatorial optimization (*e.g.* the *Traveling Salesman Problem* [2] or *Vehicle Routing Problems* [22]) because they present both practical and theoretical challenges to the researchers.

Ridesharing systems (see *e.g.*, [16, 18]) are emerging transportation models and tools where car owners can share a ride with other persons via a dedicated application (*e.g.* *avego*, *blablacar*, *carpool-world*, *carticipate*, etc.). Ridesharing systems, as public transportation systems, are valuable initiatives for the reduction of traffic congestion, CO₂ emissions and fuel expenditure.

In this article, we depart from the extensive literature dealing with centrally computed solutions (see *e.g.*, [20]) and focus on game theoretic approaches. Concerning transportation models, numerous articles on vehicle-routing games deal with cooperative games (see *e.g.*, [6, 12, 5]). However, noncooperative and competitive games are more closely related to our transportation game. For example, [13] study a competitive traveling salesmen problem in which two salespeople compete for visiting some clients earlier than their opponent. In this model, the players make their decisions in turn like in a game in perfect-information extensive form. A similar model, with possibly more than 2 salespersons, is considered in [17].

Our model of transportation differs from the aforementioned works since it is a strategic game. The literature on strategic games for routing problems can be divided in two parts, whether the players are *non-atomic* or *atomic*. In the mathematical models involving non-atomic players, there is traffic in a

network and each infinitesimal portion of this traffic is associated with an autonomous agent (see *e.g.* [10] for the notion of *Wardrop equilibrium*). In this article, we assume that the players are atomic. In comparison, an atomic player represents a non-negligible portion of the traffic.

As resource selection games, transportation games are reminiscent of *scheduling games with coordination mechanisms* [9]. In coordination mechanisms we have a scheduling policy, which imposes a priority over players in each resource. The scheduling policy may be described by some simple rule, such as *shortest (resp., largest) job first*, or may be more sophisticated (see *e.g.*, [4]). Moreover, the same or different scheduling policies can be used for the resources. The goal is to find natural coordination mechanisms that can significantly improve the resulting *price of anarchy* (see *e.g.*, [4, 7, 9]) or can ensure the existence of an equilibrium in pure strategies or the fast convergence to it (see *e.g.*, [4]). Thus, coordination mechanisms modify (or enrich) the individual cost structure, aiming at improved efficiency (or equilibrium existence).

In transportation games, we employ a fixed player priority, possibly different in each bus, to simplify the individual cost structure and to allow for an efficient best response computation. Of course, one might think of more sophisticated player priorities and bus routes, which is somewhat reminiscent to more complex scheduling policies in coordination mechanisms. Such priorities may naturally depend on the set of players in the same bus and on their distances to each other and to the destination. But, if *e.g.*, we pick the players in each bus according to the shortest route starting from the source, going through all of them, and ending up to the destination, determining such a route and the corresponding individual costs requires the solution to an NP-hard optimization problem. Keeping the player priorities fixed and independent of their partition into buses, we simplify the individual cost structure so that transportation games are amenable to theoretical analysis.

Transportation games also bear some resemblance to (non-cooperative versions of) hedonic games (see *e.g.*, [11, 8]), where the players are partitioned into coalitions and the individual cost of each player depends on the identities of other players in the same coalition (but not on the identity of the coalition). In transportation games, the players in each bus could be regarded as a coalition and the individual cost of each player depends on the identities of other players in the same bus (but in a more subtle way than in the hedonic games of *e.g.*, [15, 14]). The special case of transportation games with the same player permutation for all buses could be regarded as a hedonic game, since the individual cost of each player depends on the locations of other players in the same bus (but not on the bus itself). However, to enrich the individual cost structure of transportation games, we allow for different player permutations in the buses, which makes the individual cost of each player also depend on the bus (in addition to the locations of the players to be picked up after him). This is a significant departure from variants of hedonic games studied in the literature and a source of difficulty in establishing the existence of pure Nash equilibria.

1.3 Contribution

We conduct a theoretical analysis of the transportation game by providing answers to the following questions. Which case admits an equilibrium? Can we compute an equilibrium in polynomial time? Do the players naturally converge to an equilibrium? How good is the best (or worst) equilibrium in comparison with a social optimum?

This work only deals with pure strategy profiles (each player's choice is deterministic). A *pure Nash equilibrium* (NE) is a strategy profile (also called *state*) σ such that no player can unilaterally change her strategy and benefit [21]. A *strong equilibrium* (SE) is a refinement of the NE to group deviations. In a SE, no group of players C can jointly deviate in such a way that every member of C benefits [3]. This

article also deals with the dynamics of the transportation game. We say that the *dynamics converges* if, starting from *any* strategy profile, every series of improvement moves (better response) eventually reaches a stable state. Depending on the context (unilateral or group deviations), this stable state can be a NE or a SE.

Our results show that if all the resources have the same permutation then a SE exists and it can be computed in polynomial time (Theorem 1). But if the permutations are not identical, then there exists a simple 2-resource 3-player instance without any NE (Proposition 2). If there are 2 resources and if the distances is metric, then the dynamics converges to a NE (Theorem 2). Moreover, this equilibrium can be computed in linear time (Theorem 3). We have considered possible extensions of this positive result and there are metric 3-resource instances without any NE (Proposition 3).

Section 3 is devoted to a special metric case where distances can be 1 or 2. We provide an algorithm that computes a NE in $O(nm)$. In Section 4, we analyse the *price of anarchy* and *stability* of the transportation game under two natural social cost functions namely *egalitarian* and *Vehicle Kilometers Travelled*. These notions are worst case comparisons of the worst and best NE with a social optimum, respectively. Without metric, the price of anarchy and stability of our transportation game are unbounded, but they are bounded for metric distances. Some possible extensions and future works are discussed in the last section. Due to the space limitation, some proofs are omitted and are put in the Appendix.

2 Existence and Computation of an Equilibrium

At least two properties play an important role in the existence of an equilibrium: (i) whether the permutations of the buses are identical or not, (ii) whether the distances are metric or not.

2.1 Instances with Not Necessarily Metric Distances

Theorem 1. *If all the resources have the same permutation π , then the dynamics converges to a SE. Moreover, a SE can be built in $O(nm)$.*

Proof. Let us suppose wlog. that π is the identity permutation. Associate with each state σ a vector $\Lambda(\sigma)$ of dimension n . Define coordinate i of $\Lambda(\sigma)$ as $c_i(\sigma)$. Given a state σ , if a coalition C deviates so that σ' denotes the new state, then $\Lambda(\sigma')$ is lexicographically smaller than $\Lambda(\sigma)$. Indeed, let k be the member of C with smallest index. For $i < k$ we have $c_i(\sigma) = c_i(\sigma')$ (because of priority π , no player i is affected by the strategy of a player j when $i < k \leq j$) and $c_k(\sigma) > c_k(\sigma')$ (player k has decreased her cost). A SE σ is built as follows. Start from an empty solution and take the players in turn, from 1 to n . The current player, say i , is placed on the resource where its cost is minimum (break ties by selecting the resource with smallest index). The cost of i on resource j is $d(i, t)$ if no other player is previously assigned to j , otherwise it is $d(i, i') + c_{i'}(\sigma)$ where i' is, among the players assigned to j , the one with largest index ($i > i'$). Since $i > i'$, i is picked just before i' . The assignment of i to j does not affect the cost of i' . The resulting state is lexicographically minimal and it is built in $O(nm)$. \square

Theorem 1 cannot be extended to show the existence of a *super strong equilibrium* (SSE), even with metric distances. A SSE is a refinement of the SE where no group of players C can jointly deviate in such a way that no member of C is worst off, while at least one member is better off.

Proposition 1. *There exists a metric instance of the transportation game with $m = 2$ resources having identical permutations which does not admit any super strong equilibrium.*

distances			
	α	β	γ
t	2	100	1
α	0	0	0
β	0	0	10

states		
	1	2
σ_1	γ^*	β
σ_2	β	α
	α^*	γ

states		
	1	2
σ_3	α^*	β
	γ	β^*
σ_4	α	γ

Fig. 1. A non-metric instance of the transportation game with $m = 2$ resources and $n = 3$ players that does not admit any pure Nash equilibrium (see also Proposition 2).

The following result shows the importance of having identical permutations, if the distance is not metric.

Proposition 2. *There exists a non-metric instance of the transportation game with $m = 2$ resources and $n = 3$ players which does not admit any NE.*

Proof. Let the resources be 1 and 2 and let the players be α , β and γ . The permutations are (α, β, γ) for resource 1 and (γ, β, α) for resource 2. The distances are depicted in Fig. 1, on the left. Let us see that the instance has no pure Nash equilibrium. If β is alone on a resource (or she is the last to be picked before traveling to t) then she wants to move. In Fig. 1, on the right, the 4 remaining configurations are depicted. The player with a star on its right has incentive to switch. Hence, this instance does not admit any pure Nash equilibrium. \square

We observe that every instance of the transportation game with $n = 2$ players and $m = 2$ resources admits a SSE (actually, already in the case of NE, one player which minimizes the distance to s with highest priority in case of ties does not wish to go elsewhere). Therefore, Proposition 2 cannot be extended to the case of two players.

2.2 Instances with Metric Distances

Theorem 2. *For the transportation game with $m = 2$ resources and metric distances, better response dynamics converges to a NE.*

Proof. Suppose for the sake of contradiction that there exists a cycle in the Nash dynamics. Let $N_0 \subset N$ be the players who never change their strategy in the cycle, whereas $N_1 := N \setminus N_0 \neq \emptyset$.

Note that there is some positive integer k_j , with $j \in \{1, 2\}$, such that the k_j first players in the permutation of resource j are in N_0 and they play j . Indeed, it is a dominant strategy for the first player in the permutation of resource j to play j , because metric distances, ie., the triangle inequality imposes that the cost of that player cannot be lower (namely her distance to t).

For $j \in \{1, 2\}$, let p_j denote the player of N_1 coming first in the permutation of resource j . Let d_j be the player of N_0 who is just before p_j in the permutation of resource j . Let c_{d_j} denote the cost of d_j in the cycle which is invariant. In the cycle, if p_j plays resource j then her cost is equal to $d(p_j, d_j) + c_{d_j}$, whatever the players of $N_1 \setminus \{p_j\}$ play.

It must be $p_1 \neq p_2$, otherwise we get a contradiction with $p_1 \in N_1$ because the cost of p_1 does not depend on the strategy adopted by $N_1 \setminus \{p_1\}$.

Since the players do unilateral deviations in the cycle, there must be a state of the cycle in which p_1 and p_2 play the same strategy. Suppose wlog. that at some point p_2 profitably moves from resource 2 to resource 1, where p_1 is. The new cost of p_2 is at least her distance to p_1 plus the cost of p_1 .

$$d(p_2, d_2) + c_{d_2} > d(p_2, p_1) + d(p_1, d_1) + c_{d_1} \quad (1)$$

At some point in the cycle p_1 profitably moves to resource 2 where its cost is at least its distance to d_2 plus the cost of d_2 .

$$d(p_1, d_1) + c_{d_1} > d(p_1, d_2) + c_{d_2} \quad (2)$$

Combine inequalities (1) and (2) to get that $d(p_2, d_2) > d(p_2, p_1) + d(p_1, d_2)$, which is a violation of the triangle inequality. \square

From this proof, we know that a potential function exists. Providing an explicit potential function for metric transportation games with two resources is an open question. Another open question asks whether the Nash dynamics converges in polynomial time.

We next show how to efficiently compute one equilibrium with some central coordination (2 buses).

Theorem 3. *For $m = 2$ resources and metric distances, the transportation game has a NE that can be computed in $O(n)$.*

Proof (sketch). Let r_1 and r_2 be the two resources. We assign all players to r_1 , consider them one-by-one, as they appear in the permutation of r_2 , and let them deviate to r_2 , if this improves their cost. Moving players from r_1 to r_2 can only decrease the cost of those staying on r_1 . Hence, if a player does not move to r_2 in her turn, she does not move to r_2 later on. Using an argument similar to that in the proof of Theorem 2, we can show that after a player moves to r_2 , she never prefers to return to r_1 . Hence, this procedure reaches a pure NE in n steps. A complete proof is given in the Appendix. \square

We note that the algorithm in the proof of Theorem 3 does not build a SE. Actually, a *planar* instance with 2 resources and without any SE can be designed (see Proposition 10 in the Appendix). An instance is planar if its distances are Euclidean distances between nodes of a 2-dimensional plane (in particular, it is metric). Moreover, the following proposition shows that the existence of a NE for $m = 2$ resources (and metric distances) cannot be extended to $m = 3$.

Proposition 3. *There exists a planar instance with $m = 3$ resources which does not admit any NE.*

Proof (sketch). There are 8 players p_0, \dots, p_7 whose coordinates in the plane are $p_0 = (4, 27)$, $p_1 = (4, 1)$, $p_2 = (6, 24)$, $p_3 = (5, 6)$, $p_4 = (17, 13)$, $p_5 = (4, 29)$, $p_6 = (2, 3)$, $p_7 = (20, 21)$, and $t = (18, 6)$. There are 3 resources with permutations $(p_7 p_6 p_5 p_4 p_3 p_2 p_1 p_0)$, $(p_0 p_1 p_2 p_3 p_4 p_5 p_6 p_7)$ and $(p_3 p_2 p_6 p_7 p_0 p_5 p_4 p_1)$. This instance has 3^8 pure states. Since for the players p_7 , p_0 and p_3 selecting bus 1, bus 2 and bus 3, respectively, is a dominant strategy, the number of pure states that are candidates for an equilibrium is $3^5 = 243$. An exhaustive examination of them shows that none is a pure NE. \square

We believe that the above instance is minimal, *i.e.* any instance with at most 7 players contains a NE when $m = 3$. However, the Nash dynamics is not guaranteed to converge as indicated below.

Proposition 4. *There exists a metric instance with $m = 3$ resources and $n = 7$ players in which the Nash dynamics may cycle.*

3 Computing a Pure Nash Equilibrium for Distances 1 and 2

For the simplest case of metric distances *i.e.*, corresponding to the case that all distances are either 1 or 2, a NE exists for any number of resources and can be computed in linear time by a natural greedy algorithm. In Algorithm 1, a player is available for assignment to a resource if she is currently the first player in the resource's permutation. Among all available players, Greedy picks the player u that can be

Algorithm 1 Greedy algorithm for computing a NE

Input: set N of n players, set M of m resources, permutations π_j for all $j \in M$ (each π_j ends with \triangleleft), distance function d (we assume that $d(u, \triangleleft) = \infty$).

Output: assignment $\sigma : N \rightarrow R$ that is a PNE

for all $j \in M$ **do**
 $\text{cost}(j) \leftarrow 0$; $p_j \leftarrow t$; $n_j \leftarrow$ the first player in π_j ;

for all $v \in N$ **do**
 $\text{cost}(v) \leftarrow \infty$; $\sigma(v) \leftarrow \triangleleft$;

while $\exists j$ with $n_j \neq \triangleleft$ **do**
 $k = \arg \min_j \{\text{cost}(j) + d(p_j, n_j)\}$; $u \leftarrow n_k$;
 in case of ties, select resource k with minimum $\text{cost}(k)$
 if $\sigma(u) = \triangleleft$ **then**
 assign u to resource k and set $\sigma(u) \leftarrow k$ and $c(u) \leftarrow \text{cost}(k) + d(p_k, u)$;
 if $\sigma(u) \neq \triangleleft$ and $c(u) > \text{cost}(k) + d(p_k, u)$ **then**
 restore $\text{cost}(\sigma(u))$ and $p_{\sigma(u)}$ to their values before u 's assignment to $\sigma(u)$;
 $n_{\sigma(u)}$ becomes the first player v after u in $\pi_{\sigma(u)}$ with $c(v) > \text{cost}(\sigma(u)) + d(p_{\sigma(u)}, v)$;
 reassign u from $\sigma(u)$ to k and set $\sigma(u) \leftarrow k$ and $c(u) \leftarrow \text{cost}(k) + d(p_k, u)$;
 let n_k be the next player after u in π_k (n_k becomes \triangleleft if u is the last player);
 if $\sigma(u) = k$ **then** $p_k \leftarrow u$

return assignment σ

assigned to a resource k at a minimum cost. Ties are broken in favor of resources with minimum cost. If player u is not assigned to any resource, she is assigned to k . Otherwise, if u prefers k to her current resource, she is reassigned to k . In both cases, u is removed from the permutation of resource k and Greedy continues. We next show that Greedy terminates with a pure NE assignment if the distances are either 1 or 2.

Example 2. Consider an instance with 8 players, p_0, \dots, p_7 , and 3 resources r_1, r_2 and r_3 . The permutations are $(p_7 p_6 p_5 p_4 p_3 p_2 p_1 p_0)$, $(p_0 p_1 p_2 p_3 p_4 p_5 p_6 p_7)$ and $(p_3 p_2 p_6 p_7 p_0 p_5 p_4 p_1)$ for r_1, r_2 and r_3 , respectively. The nodes are partitioned in three sets: $\{t, p_0, p_1\}$, $\{p_2, p_3, p_4\}$ and $\{p_5, p_6, p_7\}$. The nodes in the same set are within distance 1 to each other. All other distances are 2. In Algorithm 1, at the beginning, players p_7, p_0 and p_3 are available for assignment to r_1, r_2 and r_3 , respectively. In the first iteration, p_0 is assigned to r_2 and p_1 becomes available for r_2 . In the subsequent iterations, p_7 is assigned to r_1 , p_3 to r_3 , and p_1 to r_2 (this takes place last due to the tie breaking rule). At this point, we have that $n_1 = p_6$ and $n_2 = n_3 = p_2$. Next, p_6 is assigned to r_1 , p_2 to r_3 , p_2 and p_3 are considered for and not assigned to r_2 , p_4 is assigned to r_2 , and p_5 is assigned to r_1 . From that point on, p_6, p_7, p_0, p_5, p_4 and p_1 are considered for and not assigned to r_3 , p_4, p_3, p_2, p_1 and p_0 are considered for and not assigned to r_1 , and p_5, p_6 and p_7 are not assigned to r_2 . The final assignment is $(p_7 p_6 p_5)$ to r_1 , $(p_0 p_1 p_4)$ to r_2 , and $(p_3 p_2)$ to r_3 , which is a NE.

Theorem 4. *Algorithm 1 computes a NE in $O(nm)$ time if the distances are either 1 or 2. Moreover, each player is reassigned at most once through the execution of the algorithm.*

Proof. We refer to a player u as a *candidate* for resource j if either $n_j = u$ or u appears in π_j after n_j (so u will be considered for assignment or reassignment to j in a subsequent iteration). For convenience, we let $\text{cost}^{\max} = \max_j \text{cost}(j)$ and $\text{cost}^{\min} = \min_j \text{cost}(j)$.

We use induction on the number of iterations and show that at the end of the current iteration: (i) no assigned player wants to deviate to any resource j , unless she is a candidate for j ; (ii) $\text{cost}^{\max} - \text{cost}^{\min} \leq 2$; (iii) cost^{\min} does not decrease from one iteration to the next; and (iv) if player u is

reassigned from resource $\sigma(u)$ to resource k , her cost at $\sigma(u)$ is $\text{cost}^{\max} = \text{cost}^{\min} + 2$ and her cost at k is $\text{cost}^{\min} + 1$. At the end of the algorithm, no player is a candidate for any resource. Hence, (i) implies that if Greedy terminates, the assignment σ is a pure NE.

Claims (i)-(iv) are true before the first iteration. We inductively assume that (i)-(iv) hold at the end of any iteration. To establish (i)-(iv) hold at the end of the next iteration, we distinguish between three cases: whether u is assigned for the first time to k , whether u is reassigned to k , and whether u stays with $\sigma(u)$. If u stays with $\sigma(u)$, nothing changes and (i)-(iv) remain true at the end of the current iteration.

If u is assigned or reassigned to k , u does not want to deviate at the end of the current iteration, because k minimizes u 's cost among all resources j with $n_j = u$. If u is assigned to k , other resources and assigned players are not affected, and claim (i) remains true. If u is reassigned to k , the cost of resource $\sigma(u)$ decreases. To maintain (i), we let $n_{\sigma(u)}$ be the first player after u in $\pi_{\sigma(u)}$ that wants to be assigned to $\sigma(u)$ (this may involve some backtracking in $\pi_{\sigma(u)}$). Then, if an assigned player wants to deviate to $\sigma(u)$, so as to take advantage of u 's move out of $\sigma(u)$, she has become a candidate for $\sigma(u)$. So, claim (i) holds at the end of the current iteration.

As for claims (ii)-(iv), since all distances are either 1 and 2, and due to the greedy choice of resource k and to the tie-breaking rule, k 's cost is equal to cost^{\min} . Hence, if u is assigned for the first time to k , u 's cost becomes at most $\text{cost}^{\min} + 2$. Moreover, cost^{\min} does not decrease and cost^{\max} either does not change or becomes $c(u) \leq \text{cost}^{\min} + 2$. Therefore, (ii)-(iv) hold after u 's assignment.

If u is reassigned from $\sigma(u)$ to k , the cost of u at $\sigma(u)$ is $c(u) = \text{cost}^{\max} = \text{cost}^{\min} + 2$. Furthermore, $\text{cost}(k) = \text{cost}^{\min}$ (just before u 's reassignment) and the cost of u at k (after u 's reassignment) is $\text{cost}^{\min} + 1$. These follow from the facts that $\text{cost}^{\max} - \text{cost}^{\min} \leq 2$ at the end of the previous iteration, that $c(u) > \text{cost}(k) + d(p_k, u)$ and that all distances are 1 and 2. So, after u 's reassignment, $\text{cost}(k)$ is at most cost^{\max} and $\text{cost}(\sigma(u))$ decreases by 2 and becomes cost^{\min} . These imply that (ii)-(iv) remain true after u 's reassignment.

Claims (iii) and (iv) imply that any player is reassigned at most once. Due to (iv), if a player u is reassigned from resource $\sigma(u)$ to resource k , $\text{cost}(\sigma(u))$ decreases from cost^{\max} to $\text{cost}^{\min} = \text{cost}^{\max} - 2$ and the new cost of u at k is $\text{cost}^{\min} + 1$. Thus, if u is reassigned from resource k later on, the new cost^{\min} would be $\text{cost}^{\min} - 1$, which contradicts (iii). Hence, Greedy terminates after n assignments and at most n reassignments.

Each reassignment causes a backtrack of at most n players in $\pi_{\sigma(u)}$. But only assigned players with cost $\text{cost}^{\min} + 2$ can be reassigned to $\sigma(u)$ after u moves out. So, after u is reassigned to k , we need to insert at the beginning of $\pi_{\sigma(u)}$ only assigned players that appear after u in $\pi_{\sigma(u)}$ and have cost equal to $\text{cost}^{\min} + 2$. Since there are at most m such players, the additional running time due to each reassignment is $O(m)$. So the running time of Greedy is $O(nm)$, i.e., linear in the size of the input. \square

4 The Price of Anarchy and the Price of Stability

We consider two different social functions. For a strategy profile σ , $D(\sigma)$ is the total distance travelled by the buses when they transport at least one client (for each bus we neglect the distance between s and the first client). This function reflects the environmental impact of the solution and it corresponds to the objective *Vehicle Kilometers Travelled* considered in [18]. The second function is the classical *egalitarian* social cost function $E(\sigma)$ defined as $\max_{i \in N} c_i(\sigma)$, which is also the maximum distance travelled by a single bus if the distance between s and the first client is neglected.

For every $f \in \{D, E\}$ and any given instance, σ^* denotes a state for which $f(\sigma^*)$ is minimum. The (pure) *price of anarchy* (PoA for short) is the largest ratio $f(\sigma)/f(\sigma^*)$, over all instances of the game,

where σ is a pure NE [19]. The (pure) price of stability (PoS for short) is the largest ratio $f(\sigma)/f(\sigma^*)$, over all instances of the game, where σ is the best NE with respect to f [1]. Therefore $\text{PoA} \geq \text{PoS}$.

Proposition 5. *For every $n \geq 3$, the PoS is unbounded for D or E if the distance is not metric, even if all the permutations are identical.*

Proof. If $n \leq 2$ and $m \geq 2$, $\text{PoS}=\text{PoA}= 2$ even in non-metric case. Suppose $n > m \geq 2$. Let $\varepsilon \in (0, 1)$. Player $m + 1$ is at distance $1/\varepsilon$ from every player $i \in [1; m]$. Every other distance is ε . In a NE σ , no two players i, i' satisfying $1 \leq i < i' \leq m$ are on the same bus. Thus they all have an individual cost of ε . Whichever resource player $m + 1$ selects, her individual cost is $1/\varepsilon + \varepsilon$ so $f(\sigma) \geq 1/\varepsilon$ for every $f \in \{D, E\}$. Meanwhile there exists a state such that every edge traversed by a bus has distance ε (e.g. no player $i \leq m$ is on the resource selected by player $m + 1$). When $\varepsilon \rightarrow 0$ the PoS tends to ∞ .

Suppose $3 \leq n \leq m$. Let $\varepsilon \in (0, 1)$. All node distances are ε except for the cases given below:

	t	1	2	3	\dots	n
2	3ε	ε	ε	ε	\dots	ε
3	$1/\varepsilon$	ε	$1/\varepsilon$	$1/\varepsilon$	\dots	$1/\varepsilon$

In a NE σ , players 1 and 2 are on the same resource. Player 3 cannot be picked right before player 1 so her individual cost is at least $1/\varepsilon$. It follows that $f(\sigma) \geq 1/\varepsilon$ for every $f \in \{D, E\}$. Meanwhile there exists a state such that every edge traversed by a bus has distance ε or 3ε (players 1 and 2 are on distinct resources). So $\text{PoS} \rightarrow \infty$ when $\varepsilon \rightarrow 0$. \square

4.1 Function D with metric distances

Due to Proposition 5, from now on, we assume that the distances are metric.

Lemma 1. *If d is metric, $d(x, y) \leq D(\sigma^*)$ holds for all nodes $x, y \in N \cup \{t\}$.*

Proof. If x and y are covered by the same bus in σ^* , then suppose wlog. that the bus visits x before y . Therefore, $D(\sigma^*)$ is at least the distance covered by the bus between x and y , while the latter is at least $d(x, y)$, by the triangle inequality. Hence, suppose x and y are covered by two different buses in σ^* , and denote them by b_x and b_y , respectively. Therefore, $D(\sigma^*)$ is at least the distance covered by b_x between x and t plus the distance covered by b_y between y and t . The latter is at least $d(x, y)$, by the triangle inequality. \square

Proposition 6. *If d is metric, then $D(\sigma) \leq nD(\sigma^*)$ holds for every state σ .*

Combining Lemma 1 with Proposition 6, we obtain the following upper bound on the PoA (and on the PoS). Interestingly, Proposition 7 shows that this bound is tight.

Corollary 1. *The PoA with respect to D of the transportation game on n players with metric distances and $m \geq 2$ resources is upper bounded by n .*

Proposition 7. *For any $n \geq 2$, there are metric instances of the transportation game on n players and $m \geq 2$ resources where the PoS is asymptotically n , even if all the resources have the same permutation.*

Proof. Suppose $2 \leq n \leq m$ and consider an instance where $d(i, t) = 1, \forall i \in N$. Take any δ in $(0, 1)$ and let $d(i, i') = \delta, \forall i, i' \in N$. The triangle inequality is satisfied. The permutation of each bus is the identity. There is essentially one NE in this instance: every player is the single user of a resource. The social cost is n . In an optimal state, all the players use a single resource. The optimal social cost is at most $1 + (n - 1)\delta$. Thus, for δ small but positive, the PoS tends to n .

Suppose $2 \leq m \leq n$ and consider the following instance in which k is an integer:

- $N = L \cup R$ where $L = \{\ell_{i,j} : 1 \leq i \leq k, 1 \leq j \leq m\}$ and $R = \{r_{i,j} : 1 \leq i \leq k, 1 \leq j \leq m\}$. Hence, there are $n = 2km$ players and $m \geq 2$ buses.
- $d(u, v) = 1$ if $u, v \in L$ or $u, v \in R$, $d(v, t) = a$ for some positive integer a when $v \in R$, $d(v, t) = a^2$ when $v \in L$, and finally $d(u, v) = a(1 + a)$ if $u \in L$ and $v \in R$.
- Now, let us decompose L and R into k levels (L_1, \dots, L_k) and (R_1, \dots, R_k) respectively, where $L_i = \{\ell_{i,j} : j = 1, \dots, m\}$ and $R_i = \{r_{i,j} : j = 1, \dots, m\}$.
- All the buses have the same priority π defined by $\pi(v) < \pi(u)$ if one of the following cases occurs:
 1. $v \in L_i \cup R_i$ and $u \in L_{i'} \cup R_{i'}$ with $i < i'$;
 2. $v \in L_i$ and $u \in R_i$;
 3. $v = \ell_{i,j}, u = \ell_{i,j'}$ (resp., $v = r_{i,j}, u = r_{i,j'}$) with $j < j'$.

This instance is clearly metric. We claim that a state σ is a NE iff each bus contains exactly one player of each level of L and one player of each level of R .

Obviously, the condition is sufficient because for such a state, no unilateral deviation is profitable. Conversely, by contradiction, suppose there exists a NE σ which does not satisfy the condition and let $q \geq 1$ be the first level where the condition is violated. Namely, each bus in state σ contains exactly one player of each level $L_1, \dots, L_{q-1}, R_1, \dots, R_{q-1}$. Consider the two possibilities:

- One bus, say p , does not pick any player of L_q and another one, say p' , picks at least two players, say $\ell_{q,j}$ and $\ell_{q,j'}$ where $j < j'$. By construction, $c_{\ell_{q,j'}}(\sigma) \geq c_{\ell_{q,j}}(\sigma) + d(\ell_{q,j}, \ell_{q,j'}) = c_{\ell_{q,j}}(\sigma) + 1$. Now, if player $\ell_{q,j'}$ moves from bus p' to bus p , then its cost is at most $c_{\ell_{q,j}}(\sigma)$, which contradicts the stability of σ .
- The condition holds for L_q but not for R_q . As previously, one bus, say p , has not picked any player of R_q while another bus, say p' , has picked at least two players, say $r_{q,j}$ and $r_{q,j'}$ where $j < j'$. Hence, $c_{r_{q,j'}}(\sigma) \geq c_{r_{q,j}}(\sigma) + 1 > c_{r_{q,j'}}(\sigma - r_{q,j}, p)$.

In both cases we get a contradiction, so the property holds. We deduce that $D(\sigma) = m(2k - 1)a + 2kma^2 = na^2 + a(n - m)$ iff σ is a NE. On the other hand, σ^* uses only two buses: one containing all the players of L and another containing all the players of R . We have $D(\sigma^*) = a^2 + a + (n - 2)$. When $a \rightarrow \infty$, $D(\sigma)/D(\sigma^*)$ tends to n . Finally, the result follows from Corollary 1. \square

The second instance in the proof of Proposition 7, with $a = 1$, applies to the special case where the distances are 1 or 2 (which always admits a NE by Theorem 4). So, we get an asymptotic lower bound of 2 on the PoS. An upper bound of 2 on the PoA is directly obtained from the fact that for every state σ , $D(\sigma)$ consists of $|N|$ terms which are either 1 or 2.

4.2 Function E without metric

Lemma 2. $d(x, y) \leq 2E(\sigma^*)$ holds for every pair of nodes $(x, y) \in N$, and $d(x, t) \leq E(\sigma^*)$ holds for every node $x \in N$.

Lemma 3. *In any pure Nash equilibrium, the cost of a player is at most $(2\lceil \frac{n}{m} \rceil - 1)E(\sigma^*)$.*

Proof. By contradiction, suppose there is a pure NE σ and a player i such that $c_i(\sigma) > (2\lceil \frac{n}{m} \rceil - 1)E(\sigma^*)$. Let k denote the number of players that the bus selected by i picks between i and t (this includes i). Using Lemma 2 we have $c_i(\sigma) \leq (2k - 1)E(\sigma^*)$. These bounds on $c_i(\sigma)$ give $k > \lceil \frac{n}{m} \rceil$. But if more than $\lceil \frac{n}{m} \rceil$ players use the same bus, then there must be another bus, say b , selected by less than $\frac{n}{m}$ players. Even if player i appears last in the permutation of b , her cost if she moves to b would be less than $(2\frac{n}{m} - 1)E(\sigma^*)$ (Lemma 2). We get a contradiction with the fact that σ is a NE. \square

Corollary 2. *The PoS with respect to E of the transportation game is $O(\frac{n}{m})$.*

Proof. For $n > m$, the upper bound comes from Lemma 3. The lower bound comes from the second instance in the proof of Proposition 7 in which the worst individual cost in a pure Nash equilibrium is $(2k - 1)a(1 + a) + a^2$. There is a state in which the worst individual cost is at most $2k + a^2 - 1$. For a large, the PoS tends to $2k = n/m$. \square

Proposition 8. *For the transportation game, $PoA = 2\lceil \frac{n}{m} \rceil - 1$ if $n > m$ and $PoA = 1$ if $n \leq m$.*

Proof. If $n \leq m$ then every player can be alone on her bus. By Lemma 2 (second part), the PoA is 1. Now suppose $n > m$. The upper bound comes from Lemma 3. Let us describe the lower bound.

Suppose n and m are even, especially $m = 2p$ and $n = q2p$, where q is even. Decompose N in m sets of size q as follows: $N_j = \{1^j, 2^j, \dots, q^j\}$, with $j = 1 \dots 2p$. For every player pair u, v such that $u \in N_j, v \in N_{j'}$ and $j \neq j'$ we have $d(u, v) = 2$. If u and v belong to the same set N_j then $d(u, v) = 0$. For every $u \in N$, $d(u, t) = 1$. Thus the triangle inequality holds. If $j \in [1, 2p]$ is odd then the first players of the permutation of resource j are $(1^j, 1^{j+1}, 3^j, 3^{j+1}, \dots, q - 1^j, q - 1^{j+1})$ followed by the rest of the players in an arbitrary order. If $j \in [1, 2p]$ is even then the first players of the permutation of resource j is $(2^j, 2^{j-1}, 4^j, 4^{j-1}, \dots, q^j, q^{j-1})$ followed by the rest of the players in an arbitrary order.

Consider the state σ where each resource $j \in [1, 2p]$ is selected by the first $2q$ players in its permutation. This state is a pure Nash equilibrium because every player has cost at most $4q - 1$ and deviating induces a cost of $4q + 1$. In an optimal state the players of N_j are on resource j ; the maximum cost is 1. Therefore the PoA is at least $4q - 1 = 2n/m - 1$. \square

We can also bound the PoA according to the parameters $d_{\min} = \min_{x \neq y} d(x, y)$ and $d_{\max} = \max d(x, y)$. As an immediate corollary, we obtain that $PoS = PoA = 2$ for distances 1 and 2.

Proposition 9. $PoS = PoA = d_{\max}/d_{\min}$ for the transportation game.

5 Future Directions

In this work we supposed that the route of the buses are prescribed by a permutation that is independent of the current state. There is an interesting challenge of proposing different ways to define the route of the buses. This modification would induce a different structure of the individual costs and possibly provide better PoA and PoS, under the constraint that a pure equilibrium exists. This challenge is similar to the search of coordination mechanisms in scheduling games.

In transportation problems, it is important to predict the situation so it would be interesting to identify the cases where the equilibrium is unique. In the future, the model of transportation can be extended in several natural aspects. Each bus may have a capacity, its own speed and dedicated roads. The players may have different sizes (*e.g.* a player is a group of persons).

References

1. Elliot Anshelevich, Anirban Dasgupta, Jon M. Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. In *45th Symposium on Foundations of Computer Science (FOCS 2004), 17-19 October 2004, Rome, Italy, Proceedings*, pages 295–304. IEEE Computer Society, 2004.
2. David L. Applegate, Robert E. Bixby, Vasek Chvatal, and William J. Cook. *The Traveling Salesman Problem: A Computational Study (Princeton Series in Applied Mathematics)*. Princeton University Press, Princeton, NJ, USA, 2007.
3. Robert J. Aumann. Acceptable points in general cooperative n-person games. In Albert W. Tucker and Robert D. Luce, editors, *Contribution to the Theory of Games*, volume IV of *Annals of Mathematics Studies*, 40, pages 287–324. Princeton Univ. Press, 1959.
4. Yossi Azar, Kamal Jain, and Vahab S. Mirrokni. (almost) optimal coordination mechanisms for unrelated machine scheduling. In Shang-Hua Teng, editor, *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2008, San Francisco, California, USA, January 20-22, 2008*, pages 323–332. SIAM, 2008.
5. Filippo Bistaffa, Alessandro Farinelli, and Sarvapali D. Ramchurn. Sharing rides with friends: A coalition formation algorithm for ridesharing. In Blai Bonet and Sven Koenig, editors, *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA.*, pages 608–614. AAAI Press, 2015.
6. Peter Borm, Herbert Hamers, and Ruud Hendrickx. Operations research games: A survey. *Top*, 9(2):139–199, 2001.
7. I. Caragiannis. Efficient coordination mechanisms for unrelated machine scheduling. *Algorithmica*, 66(3):512–540, 2013.
8. Katarína Cechlárová. Stable partition problem. In Ming-Yang Kao, editor, *Encyclopedia of Algorithms*. Springer, 2008.
9. George Christodoulou, Elias Koutsoupias, and Akash Nanavati. Coordination mechanisms. *Theor. Comput. Sci.*, 410(36):3327–3336, 2009.
10. J. R. Correa and N. E. Stier-Moses. Wardrop equilibria. *Wiley Encyclopedia of Operations Research and Management Science.*, 2011.
11. J. H. Drèze and J. Greenberg. Hedonic coalitions: Optimality and stability. *Econometrica*, 48(4):pp. 987–1003, 1980.
12. S. Engevall, M. Lundgren, and P. Värbrand. The heterogeneous vehicle - routing game. *Transportation Science*, 38(1):71–85, 2004.
13. Sándor P. Fekete, Rudolf Fleischer, Aviezri Fraenkel, and Matthias Schmitt. Traveling salesmen in the presence of competition. *Theoretical Computer Science*, 313(3):377 – 392, 2004. Algorithmic Combinatorial Game Theory.
14. Moran Feldman, Liane Lewin-Eytan, and Joseph Naor. Hedonic clustering games. In Guy E. Blelloch and Maurice Herlihy, editors, *24th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '12, Pittsburgh, PA, USA, June 25-27, 2012*, pages 267–276. ACM, 2012.
15. Martin Gairing and Rahul Savani. Computing stable outcomes in hedonic games. In Spyros C. Kontogiannis, Elias Koutsoupias, and Paul G. Spirakis, editors, *Algorithmic Game Theory - Third International Symposium, SAGT 2010, Athens, Greece, October 18-20, 2010. Proceedings*, volume 6386 of *Lecture Notes in Computer Science*, pages 174–185. Springer, 2010.
16. Ece Kamar and Eric Horvitz. Collaboration and shared plans in the open world: Studies of ridesharing. In Craig Boutilier, editor, *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence, Pasadena, California, USA, July 11-17, 2009*, page 187, 2009.
17. Graham Kendall and Jiawei Li. Competitive travelling salesmen problem: A hyper-heuristic approach. *Journal of the Operational Research Society*, 64:pp. 208–216, 2013.
18. Alexander Kleiner, Bernhard Nebel, and Vittorio A. Ziparo. A mechanism for dynamic ride sharing based on parallel auctions. In Toby Walsh, editor, *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 266–272. IJCAI/AAAI, 2011.
19. Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. In Christoph Meinel and Sophie Tison, editors, *STACS 99, 16th Annual Symposium on Theoretical Aspects of Computer Science, Trier, Germany, March 4-6, 1999, Proceedings*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413. Springer, 1999.
20. Hoong Chuin Lau, Lucas Agussurja, Shih-Fen Cheng, and Pang Jin Tan. A multi-objective memetic algorithm for vehicle resource allocation in sustainable transportation planning. In Francesca Rossi, editor, *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*. IJCAI/AAAI, 2013.
21. John Nash. Non-cooperative Games. *The Annals of Mathematics*, 54(2):286–295, 1951.
22. Paolo Toth and Daniele Vigo, editors. *The Vehicle Routing Problem*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2001.

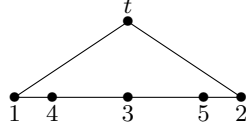


Fig. 2. The instance used in the proof of Proposition 1

A Appendix

Proposition 1. *There exists a metric instance of the transportation game with $m = 2$ resources having identical permutations which does not admit any super strong equilibrium.*

Proof. Consider the instance depicted in Fig. 2 with $m = 2$ resources available. We have $d(1, t) = d(2, t)$, $d(3, 1) = d(3, 2)$, $d(4, 1) = d(5, 2)$ and $d(3, t) > 0$. The permutation of each resource is the identity. In order to break ties, resource 1 is always chosen if the two resources induce the same cost. A super strong equilibrium σ must be a Nash equilibrium so we can suppose that $\sigma_1 = 1$ and $\sigma_2 = 2$. Now it suffices to consider all possible strategies for players 3, 4 and 5. Actually, for every profile, there exists within $\{3, 4, 5\}$ a non-empty coalition of players who can deviate so that the cost of every member of the coalition does not increase and the individual cost of at least one player strictly decreases. E.g., consider the state $\sigma = (1, 2, 2, 1, 1)$. We have $c_3(\sigma) = d(3, 2) + d(2, t)$, $c_4(\sigma) = d(4, 1) + d(1, t)$ and $c_5(\sigma) = d(5, 4) + d(4, 1) + d(1, t)$. If players 3 and 5 switch, then $\sigma' = (1, 2, 1, 1, 2)$. We have $c_3(\sigma') = d(3, 1) + d(1, t) = c_3(\sigma)$ and $c_5(\sigma') = d(5, 2) + d(2, t) < c_5(\sigma)$. \square

Theorem 3. *For $m = 2$ resources and metric distances, the transportation game has a NE that can be computed in $O(n)$.*

Proof. Let r_1 and r_2 be the two resources. We let wlog. the permutation of r_1 be the identity permutation and let π_2 be the permutation of r_2 . We assign all players to r_1 , consider them one-by-one, as they appear in π_2 , and let them deviate to r_2 if this improves their cost. More formally, for $i = 1, \dots, n$, player $j = \pi_2(i)$ computes her cost $c_j(r_1)$ in r_1 and $c_j(r_2)$ in r_2 , in the current configuration. Player j stays in r_1 , if $c_j(r_1) \leq c_j(r_2)$, and she moves to r_2 , otherwise.

Moving players from r_1 to r_2 can only decrease the completion time of those staying on r_1 . Hence, if a player does not move to r_2 in her turn, she does not move to r_2 later on. We next show that after a player moves to r_2 , she never prefers to return to r_1 . Therefore, the procedure above reaches a pure NE in n steps.

For sake of contradiction, let j be a player on r_2 that can improve her cost by returning to r_1 and let k be the first player whose move from r_1 to r_2 makes j prefer r_1 to r_2 . For convenience, we let $c_j(r_1)$ and $c_j(r_2)$ (resp. $c'_j(r_1)$ and $c'_j(r_2)$) be j 's cost on r_1 and r_2 before (resp. just after) k moves from r_1 to r_2 . We also use similar notation for k . By hypothesis, $c'_j(r_1) < c'_j(r_2)$ and $c_k(r_1) > c_k(r_2)$. Since the players move to r_2 in the order they appear in π_2 , k 's move to r_2 does not increase j 's completion time on r_2 , i.e., $c'_j(r_2) = c_j(r_2)$. Moreover, by the triangle inequality,

$$c_k(r_2) \geq d(k, j) + c_j(r_2) = d(k, j) + c'_j(r_2) \quad (3)$$

since r_2 has to move from k to j and then from j to t . Since j prefers r_1 to r_2 after k 's move, k is picked after j in r_1 . Let $\ell < k$ be the node picked immediately after k in r_1 in the configuration just before k 's move from r_1 to r_2 . Then, the triangle inequality implies that

$$c_k(r_1) \leq d(k, j) + d(j, \ell) + c_\ell(r_1) \leq d(k, j) + c'_j(r_1) \quad (4)$$

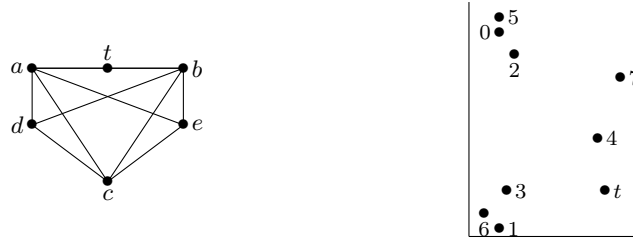


Fig. 3. On the left, an illustration of the instance used in the proof of Proposition 10. On the right, an illustration of the instance used in the proof of Proposition 3.

The second inequality holds because j is picked before k , and thus before ℓ , in r_1 . Hence, if j moves back to r_1 , r_1 has to move from j to ℓ (possibly passing through other nodes) and then from ℓ to t . Hence, $c'_j(r_1) \geq d(j, \ell) + c_\ell(r_1)$.

Combining (3) and (4) with $c_k(r_1) > c_k(r_2)$, since k moves from r_1 to r_2 , we obtain that

$$d(k, j) + c'_j(r_2) \leq c_k(r_2) < c_k(r_1) \leq d(k, j) + c'_j(r_1)$$

Therefore, $c'_j(r_1) > c'_j(r_2)$, after k moves from r_1 to r_2 , which contradicts the hypothesis that j wants to return to r_1 . \square

Proposition 10. *For $m = 2$ resources, there exists a planar instance without any strong equilibrium.*

Proof. Consider the planar instance depicted in Fig. 3. There are 5 players a, b, c, d and e . The permutations on resources 1 and 2 are (b, d, e, c, a) and (a, e, c, d, b) , respectively. This instance admits two pure Nash equilibria which are not strong equilibria.

- In the first Nash equilibrium b, d and c are on resource 1 whereas a and e are on resource 2. Therefore d and e can simultaneously switch their strategies and benefit.
- In the second Nash equilibrium b, d are on resource 1 whereas a, e and c are on resource 2. Therefore d, e and c can simultaneously switch their strategies and benefit. \square

Proposition 3. *There exists a planar instance with $m = 3$ resources which does not admit any NE.*

Proof. There are 8 players denoted by p_0, \dots, p_7 . Their coordinates in the plane are given in the next table.

p_0	p_1	p_2	p_3	p_4	p_5
(4, 27)	(4, 1)	(6, 24)	(5, 6)	(17, 13)	(4, 29)
p_6	p_7	t			
(2, 3)	(20, 21)	(18, 6)			

There are 3 resources with the following permutations.

	high priority \rightarrow low priority
bus 1	$p_7 \ p_6 \ p_5 \ p_4 \ p_3 \ p_2 \ p_1 \ p_0$
bus 2	$p_0 \ p_1 \ p_2 \ p_3 \ p_4 \ p_5 \ p_6 \ p_7$
bus 3	$p_3 \ p_2 \ p_6 \ p_7 \ p_0 \ p_5 \ p_4 \ p_1$

This instance has 3^8 pure states. Since for the players p_7, p_0 and p_3 selecting bus 1, bus 2 and bus 3, respectively, is a dominant strategy, the number of pure states that are candidates for an equilibrium is $3^5 = 243$. An exhaustive examination of them shows that none is a pure NE (see also Fig. 3, on the right, for an illustration of the instance). \square

Proposition 4. *There exists a metric instance with $m = 3$ resources and $n = 7$ players in which the Nash dynamics may cycle.*

Proof. The metric space and the permutation of each resource are depicted on Figure 4. The priority is given from bottom to top, i.e. β is visited last on resource 1. The cycle consists of 10 states which are given in Table 1. In state σ_1 , player a can profitably deviate and play resource 1, resulting in state σ_2 . The next state is σ_3 where a star indicates which player moves, and so on. Finally, we can move from σ_{10} back to σ_1 . \square

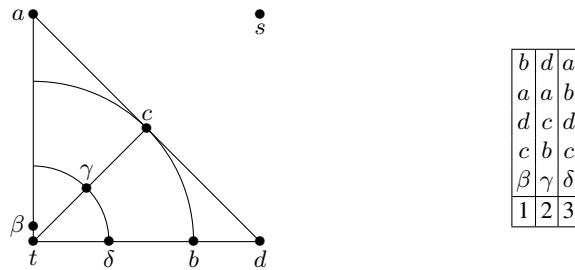


Fig. 4. The metric instance (on the left) and the permutation of each resource (on the right) used in the proof of Proposition 4

	1	2	3
σ_1	b	c	d
σ_2	a	c	d
σ_3	a	b	d
σ_4	a	b	c

	1	2	3
σ_5	d	b	c^*
σ_6	d		
σ_7	d		b

	1	2	3
σ_8	d^*	a	b
σ_9	c^*	a	d
σ_{10}		a	b^*

Table 1. The 10 states of a cycle. For each state we suppose that β plays 1, γ plays 2 and δ plays 3. The star indicates which player profitably deviates.

Proposition 6. *if d satisfies the triangle inequality, then $D(\sigma) \leq nD(\sigma^*)$ holds for every pure strategy profile σ .*

Proof. For $j \in M$ and a strategy profile σ , let $\{j_1, \dots, j_{n_j}\}$ be the players picked in this order by bus j . Thus, $D(\sigma) = \sum_{j \leq m} \sum_{i=1}^{n_j} d(j_i, j_{i+1})$ where $j_{n_j+1} = t$ for all bus. $D(\sigma)$ consists of n terms and each one is at most $D(\sigma^*)$ by Lemma 1. \square

Lemma 2. *$d(x, y) \leq 2E(\sigma^*)$ holds for every pair of nodes $(x, y) \in N$, and $d(x, t) \leq E(\sigma^*)$ holds for every node $x \in N$.*

Proof. For the first point if x and y are covered by the same bus in σ^* then suppose wlog. that the bus visits x before y . Therefore $2E(\sigma^*) \geq E(\sigma^*) \geq c_y(\sigma^*) + d(x, y) \geq d(x, y)$. If x and y are covered by two different buses in σ^* then $d(x, y) \leq d(x, t) + d(y, t) \leq c_x(\sigma^*) + c_y(\sigma^*) \leq 2E(\sigma^*)$. For the second point we have $E(\sigma^*) \geq c_x(\sigma^*) \geq d(x, t)$. \square

Proposition 9. *PoS = PoA = d_{\max}/d_{\min} for the transportation game.*

Proof. Let σ and σ^* be a pure Nash equilibrium and a socially optimal state, respectively. Let ℓ^* denote the maximum number of players on the same resource in σ^* . Thus $E(\sigma^*) \geq \ell^* d_{\min}$. By contradiction, suppose $E(\sigma) > d_{\max} E(\sigma^*) \geq d_{\max} \ell^*$. Any player, say i , who has an individual cost of $E(\sigma)$ in σ shares her bus with at least ℓ^* other players. There must be a bus b , such that $\sigma_i \neq b$, selected by at most $\ell^* - 1$ players in σ , otherwise there are strictly more players in σ than in σ^* . If player i moves to bus b then her individual cost is at most $d_{\max} \ell^*$. It contradicts the fact that σ is a Nash equilibrium, so $d_{\min} E(\sigma) \leq d_{\max} E(\sigma^*)$.

The second instance given in the proof of Proposition 7 provides a lower bound of $(4k - 1)/2k$ for the price of stability for $d_{\max} = 2d_{\min}$. \square