Network Topology and Equilibrium Existence in Weighted Network Congestion Games

Igal Milchtaich*

International Journal of game theory 44 (2015), 515–541

Abstract. Every finite game can be represented as a weighted network congestion game on some undirected two-terminal network. The network topology may reflect certain properties of the game. This paper solves the *topological equilibrium-existence problem* of identifying all networks on which *every* weighted network congestion game has a pure-strategy equilibrium. *JEL Classification:* C72

Keywords: Network games, congestion games, existence of pure-strategy equilibrium, finite improvement property

1 Introduction

A weighted network congestion game is played on a network, where each player has to choose a route connecting the players' common origin and destination vertices. The players' alternatives may differ, however, since not all of them are necessarily allowed to use all edges. Players may differ also in their weights, which quantify their contributions to congestion at the edges belonging to their routes. As congestion increases, an edge's cost weakly increases or the gain from using it weakly decreases.

Games of this kind may naturally be used to model negative externalities due to limited network resources. Costs may represent, for example, travel or service times and weights may represent the agents' congestion impacts or their demands (assuming that these cannot be split among multiple routes). However, somewhat surprisingly, weighted network congestion games may also serve as concrete representations of arbitrary normal-form games: *every* finite game can be represented as (in other words, it is isomorphic to) such a network game (Milchtaich 2013).

This representation result raises the question of what properties of the represented game can be inferred from the representation, in particular, from the network used. A particularly interesting property is the existence of at least one pure-strategy Nash equilibrium. A network has the *(equilibrium-) existence property* if every weighted network congestion game on it has a pure-strategy equilibrium, which implies the same for every finite game that can be represented as such a network game. The last implication may be relevant even if the finite game is completely specified, so that, in principle, the existence of pure-strategy equilibrium can be determined by exhaustive search. This is because doing so may take a long time even if the number of players is only moderately large. In fact, the equilibrium-existence decision problem is NP-complete even for finite games where no player has more

* Department of Economics, Bar-Ilan University, Ramat Gan 5290002, Israel
igal.milchtaich@biu.ac.il https://sites.google.com/view/milchtaich/

than two strategies, two possible payoffs and two other players who may affect his payoff (Fischer et al. 2006).

This paper solves the *topological equilibrium-existence problem*, which is the identification of all (undirected) two-terminal networks with the existence property. The problem was first raised by Libman and Orda (2001), who were also the first to give an example of a (two-player) weighted network congestion game without a pure-strategy equilibrium. (For other versions of the same example, see Fotakis et al. 2005, Goemans et al. 2005, and Figure 5 below). The solution builds on partial results obtained in Milchtaich (2006b). However, these results pertain to a somewhat different definition of weighted network congestion games, which in particular assumes that all players are allowed to use all edges and thus renders these network games incapable of representing finite games where different players have different numbers of strategies. This class of weighted network congestion games with public edges and several other related models are briefly discussed in Section 4.

One of the related models concerns (unweighted) network congestion games with player-specific costs. In these games, players have identical weights, but on the other hand, the edges' cost functions are not the same for all players. Like weighted network congestion games, these network games are capable of representing all finite games (Milchtaich 2013).[1] However, for them, the topological equilibrium-existence problem is still open.

2 Preliminaries

2.1 Game theory

A finite (noncooperative) game Γ has a finite number n of players, numbered from 1 to n. Each player i has a finite set of strategies S_i^Γ and a payoff function h_i^Γ that specifies i's payoff for each strategy profile $(s_1, s_2, \dots, s_n)$. A strategy profile is a pure-strategy (Nash) equilibrium if none of the players can increase his payoff by unilaterally switching to any other strategy.

Two finite games Γ and Γ' with the same number n of players are *isomorphic* (Monderer and Shapley 1996; Milchtaich 2013) if it is possible to

(i) renumber[2] the players in Γ' and

(ii) find a bijection $\phi_i: S_i^\Gamma \to S_i^{\Gamma'}$ from the strategy set of each player i in Γ to that of player i (according to the new numbering) in Γ'

such that for every strategy profile $(s_1, s_2, \dots, s_n)$ in Γ

$$h_i^\Gamma(s_1, s_2, \dots, s_n) = h_i^{\Gamma'}(\phi_1(s_1), \phi_2(s_2), \dots, \phi_n(s_n)), \qquad i = 1,2, \dots, n. \tag{1}$$

[1] The existence of representations of a similar kind was first indicated by Monderer (2007).

[2] Renumbering effectively pairs each player i in Γ with a specific player in Γ', namely, the one (re)assigned the same number i. Therefore, it could alternatively be *defined* as a bijection between the two sets of players, that is, a one-to-one mapping from the player set in Γ onto that in Γ'.

Essentially, isomorphic games are just alternative representations of a single game. In particular, Γ has a pure-strategy equilibrium $(s_1, s_2, \dots, s_n)$ if and only if Γ' has such an equilibrium (namely, $(\phi_1(s_1), \phi_2(s_2), \dots, \phi_n(s_n))$).

Two games Γ and Γ' with identical sets of players and respective strategy sets are *similar* if, for every strategy profile $(s_1, s_2, \dots, s_n)$, the change in the payoff of a player i who unilaterally switches to another strategy s_i' is the same in both games:

$$h_i^{\Gamma}(s_1, s_2, \dots, s_i', \dots, s_n) - h_i^{\Gamma}(s_1, s_2, \dots, s_i, \dots, s_n) = h_i^{\Gamma'}(s_1, s_2, \dots, s_i', \dots, s_n) - h_i^{\Gamma'}(s_1, s_2, \dots, s_i, \dots, s_n).$$

Equivalently, for each player i, the difference $h_i^{\Gamma} - h_i^{\Gamma'}$ between player i's payoffs in the two games is unaffected by changing only i's own strategy, and can therefore be expressed as a function of the other players' strategies. Similarity implies that the two games are *best-response equivalent* (Monderer and Shapley 1996; Morris and Ui 2004), that is, a player's strategy is a best response to the other players' strategies in one game if and only if this is so in the other game. In particular, similar games have identical sets of pure-strategy equilibria.

A game Γ is an *exact potential game* (Monderer and Shapley 1996) if it is similar to some game Γ' in which all players have the same payoff function; that function P is said to be an *exact potential* for Γ. Note that this concept is a cardinal one: an increasing transformation of payoffs does not generally transform an exact potential game into another such game. An ordinal generalization of exact potential is *generalized ordinal potential* (Monderer and Shapley 1996), or simply *potential*, which is defined as a real-valued function over strategy profiles that (strictly) increases whenever a single player changes his strategy and increases his payoff as a result. Clearly, if a potential exists, then its (even "local") maximum points are equilibria. However, the existence of a potential in a finite game implies more than the existence of equilibrium. It is equivalent to the *finite improvement property* (Monderer and Shapley 1996): every improvement path (which is a finite sequence of strategy profiles where each profile differs from the preceding one only in the strategy of a single player, whose payoff increases as a result of the change) is finite. In other words, the game has no improvement cycles (which are finite improvement paths that start and terminate with the same profile). A potential does not necessarily exist in finite games that only possess the weaker *finite best-(reply) improvement property* (Milchtaich 1996). This property differs from the finite improvement property in only requiring finiteness of best-(reply) improvement paths (where in each step the new strategy is also a best response for the moving player) or equivalently nonexistence of best-improvement cycles.

The *superposition* of a finite number m of games with identical sets of players is the game with the same set of players where each player has to choose one of his strategies in each of the m games and his payoff is the sum of the resultant m payoffs (von Neumann and Morgenstern 1953). Thus, the m games are played simultaneously but independently. It is easy to see that a strategy profile in the superposition of m games is an equilibrium if and only if it induces (by projection) an equilibrium in each of the constituent m games.

2.2 Graph theory

An *undirected multigraph* consists of a finite set of vertices and a finite set of edges. Each edge e joins two distinct vertices, which are referred to as the *end vertices* of e. Thus, loops are not allowed but more than one edge can join two vertices. An edge e and a vertex v are *incident* with each other if v is an end vertex of e. The *degree* of a vertex is the number of edges incident with it. A (simple) *path* of length m is an alternating sequence of vertices and edges $v_0 e_1 v_1 \cdots v_{m-1} e_m v_m$, beginning and ending with vertices, in which each edge is incident with the two vertices immediately preceding and following it and all the vertices (and necessarily all the edges) are distinct. If the first and last vertices are clear from the context, the path may be written more simply as $e_1 e_2 \cdots e_m$. Every path traverses each of its edges e in a particular *direction*: from the end vertex that immediately precedes e in the path to the vertex that immediately follows it.

A *two-terminal network*, or simply *network*, G is an undirected multigraph with a pair of distinguished *terminal* vertices, a vertex o called the *origin* and another one d called the *destination*, such that each of the vertices and edges in the multigraph belongs to at least one path that begins with o and ends with d. Such a path is called a *route* in G. A route may itself be viewed as a network. Indeed, it is an example of a *sub-network* of G, that is, a network that can be obtained from G by deleting some of its edges and non-terminal vertices.

Two networks are *isomorphic* if there is a one-to-one correspondence between their sets of vertices, and another such correspondence between the sets of edges, such that the incidence relation is preserved and the origin and destination in one network are paired with the origin and destination, respectively, in the other network. Isomorphic networks may be, and they normally are, viewed as identical: two copies of the same network.

It can be shown that performing any of the following operations on a network G creates a new network with the same terminal vertices (see Figure 1):

(a) *Subdivision of an edge*: its replacement by two edges with a single common end vertex.
(b) *Addition of a new edge* joining two existing vertices.
(c) *Subdivision of a terminal vertex*: addition of a new edge e joining o or d with a new vertex v, followed by the replacement of the terminal vertex by v as the end vertex in two or more edges originally incident with the former.

A network G is *embedded in the wide sense*[3] in a network G' if the latter can be obtained from the former by applying the above operations any number of times in any order. Every sub-network of a network is embedded in the wide sense in it (Milchtaich 2005).

[3] Embedding in the wide sense, which was introduced in Milchtaich (2005), is more inclusive than the narrower notion of embedding used in Milchtaich (2006a). The difference is that, in the latter, the only kind of terminal subdivision (operation (c) above) allowed is *terminal extension*, in which *all* the edges originally incident with the terminal vertex become incident with the new vertex v instead. Whereas embedding in the wide sense roughly corresponds to the notion of a minor of a graph, embedding in the narrower sense corresponds to a topological minor (see Diestel 2005).

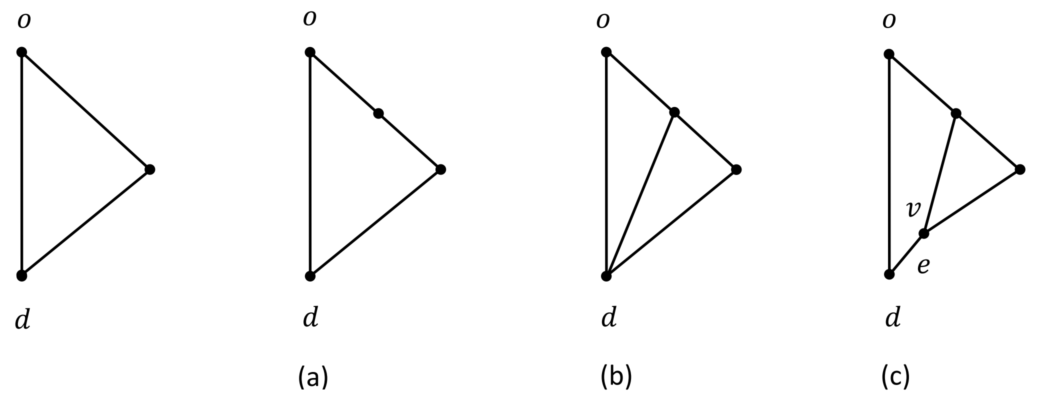

Figure 1. Embedding. The left-most network is embedded in the wide sense in each of the other three, which are obtained from it by (a) subdividing an edge, (b) adding a new edge, and, finally, (c) subdividing the destination.

Two networks G and G' are *homeomorphic* if they can be obtained from the same network by successive subdivision of edges, in other words, if each of them can be obtained from the other by the insertion and removal of non-terminal vertices of degree two. While technically distinct, homeomorphic networks are topologically similar and, from the perspective of network congestion games (of the kinds considered below), practically identical. In Figure 1, the two left networks are homeomorphic to one another and to the network with only two (terminal) vertices and two edges.

A network G may be connected with another network G', which does not share any of its edges and vertices, *in series* or *in parallel*. The sets of vertices and edges in the resultant network are the unions of the corresponding sets in G and G', except that: for a connection in series, the destination in G and the origin in G' are identified, and become a single non-terminal vertex; and for a connection in parallel, the two origin vertices as well as the two destination vertices are identified, and become a single pair of terminal vertices. For example, connecting a network that only has one edge in series with the left-most network in Figure 1, and then connecting the resultant network in parallel with a second single-edge one, gives the network in (b). The connection of an arbitrary number of networks in series or in parallel is defined recursively. Each of the connected networks is embedded in the wide sense in the resultant one.

2.3 Network congestion games

A *weighted network congestion game* on a (two-terminal[4]) network G is a finite, n-player game that is defined as follows. First, each edge e in G is assigned a nondecreasing *cost function*[5] $c_e: (0, \infty) \to (-\infty, \infty)$, an allowable direction, which must be that in which some route in G traverses e, and a (possibly, empty) set of allowable users. An edge is *public* or *private* if it is allowable to all players or to one player only, respectively. It is required that

[4] The assumption of a single origin–destination pair may be viewed as a normalization. Any weighted network congestion game on a *multi-commodity network*, which has multiple origin–destination pairs, may also be viewed as a game with a single such pair. In that game, each of the two terminal vertices is incident with a single allowable edge (see below) for each player, which joins it with the player's corresponding terminal vertex in the original game.

[5] The definition of cost function allows for negative costs, which may be interpreted as (net) gains from using the edge. However, negative costs do not play any role in Section 3, where all the results would hold also with the more restrictive definition that only allows nonnegative cost functions, $c_e: (0, \infty) \to [0, \infty)$.

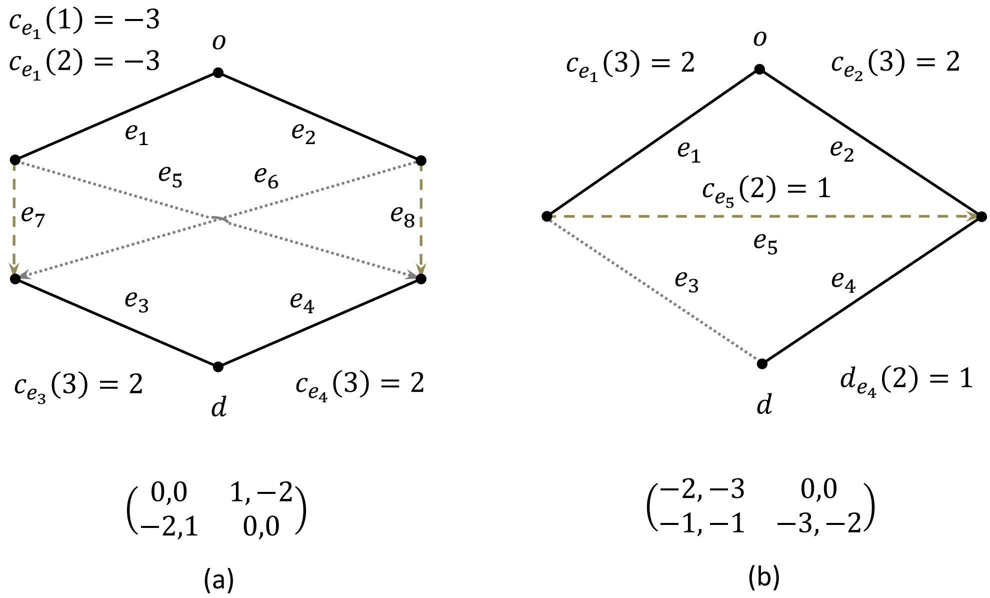

Figure 2. A two-player weighted network congestion game (a) and a two-player weighted network congestion game in the wide sense (b). In both games, the players' weights are $w_1 = 1$ and $w_2 = 2$. Dotted, dashed and solid edges are allowable to player 1, player 2 and both players, respectively. The allowable directions are indicated where needed. All relevant costs other than those specified are zero. A player's payoff is the negative of his total cost. The games' normal (or strategic) forms are shown at the bottom.

each player i has at least one *allowable route*, that is, a route in G that includes only edges that i is allowed to use and traverses them in the allowable direction. The collection of all such routes is the player's strategy set S_i. Second, a *weight* $w_i > 0$ is specified for each player i, which represents the player's congestion impact and is also (weakly) connected with the cardinality of his strategy set: For all i and j with $w_i < w_j$, $|S_i| \geq |S_j|$.[6] The total weight f_e of the players whose chosen route includes an edge e is the *flow* (or load) on e. The *cost* of e for each of its users is $c_e(f_e)$. A player's payoff in the game is the negative of the total cost of the edges in his route.

A weighted network congestion game is referred to as an *unweighted* network congestion game if the players' weights are all identical and equal to 1. The equality of the weights entails, in particular, that the cost of an edge is not affected by the identities of its users but only by their number. A generalization that allows for a dependence of the cost for a user on his own identity is (unweighted) *network congestion game with player-specific costs*. In such a game, each edge e is associated with a nondecreasing cost function $c_{ie}: (0, \infty) \to (-\infty, \infty)$ for each player i, and its cost for that player is $c_{ie}(f_e)$, where (the flow) f_e is the total number of players using e.

The very specific definitions of network congestion games make them appear special. However, as the following representation theorem (Milchtaich 2013) shows, in a very fundamental sense, this is not so.

[6] The cardinality assumption is used in the proof of Lemma 5. Whether or not it can be dispensed with I do not know. In one, important sense, the assumption is not overly restrictive. The proof of the representation theorem (Theorem 1 below) only uses network games that satisfy the assumption, which means that the theorem's results hold with as well as without it (Milchtaich 2013). Note that the cardinality assumption trivially holds if all players have the same number of strategies, or if the allowable users of each edge are those whose weight does not exceed a certain threshold.

Theorem 1. Every finite game Γ is isomorphic both to a weighted network congestion game Γ' and to an (unweighted) network congestion game with player-specific costs Γ''. Γ is isomorphic to an unweighted network congestion game[7] if and only if it is an exact potential game.

An immediate corollary of the next lemma is that a representation of a finite game Γ as a particular variety of network congestion game is never unique.

Lemma 1. If a network G is homeomorphic to a network G' or is embedded in the wide sense in it, then every weighted network congestion game, unweighted network congestion game, or network congestion game with player-specific costs on G is isomorphic to a game of the same kind on G'.

Proof. By the definitions of homeomorphism and embedding in the wide sense, it suffices to consider the special case in which either G' is obtained from G by one of the three operations defining embedding in the wide sense (Figure 1) or G is obtained from G' by the subdivision of an edge. Given a network congestion game Γ on G, it has to be shown that an isomorphic game of the same kind exists on G'. Such a game Γ' can be obtained by "extending" Γ to G', that is, assigning a cost function (or cost functions, if Γ is a network congestion game with player-specific costs), an allowable direction and a set of allowable users to each of the (one or two) edges in G' that are not in G. The assignments are as follows.

(a) If the operation connecting G and G' is the subdivision of an edge e in G, each of the two "halves" of e is assigned half its cost and inherits its allowable direction and set of allowable users.
(b) If the operation is the addition of a new edge e to G, no player is allowed to use e.
(c) If the operation is the subdivision of a terminal vertex in G, which creates the new edge e (and a new vertex v), all players are allowed to use e at zero cost. Since e is incident with a terminal vertex, is has only one possible allowable direction.
(d) If the operation is the subdivision of an edge e' in G', e' is assigned the sum of its two halves' costs in Γ and the allowable direction of either of them, and its allowable users are all players who, in Γ, have an allowable route that includes the two halves.

Consider, for each player i, the following function ϕ_i from the set of i's allowable routes in G to that in G'. If s_i includes an edge e as in (a), $\phi_i(s_i)$ is the route obtained from s_i by replacing e with its two halves and their common end vertex. If s_i includes the two halves of an edge e' as in (d) and their common end vertex, $\phi_i(s_i)$ is obtained by replacing these with e'. If s_i includes an edge that in G is incident with a particular terminal vertex (o or d) but in G' the latter is replaced with a different end vertex v, which is shared with an edge e as in (c), $\phi_i(s_i)$ is obtained from s_i by inserting e and v next to the terminal vertex. For any other s_i, $\phi_i(s_i) = s_i$. It is not difficult to see that ϕ_i is a bijection and that the identity (1) holds. Therefore, the games Γ and Γ' are isomorphic. ■

[7] This condition can be expressed as the requirement that $\Gamma' = \Gamma''$.

It follows from Lemma 1 that the collection of all networks on which a finite game Γ is representable is completely determined by the collection's *minimal* elements, i.e., those in which no other element is embedded in the wide sense. There may be more than one such minimal network. For example, it is easy to see that if Γ is representable (minimally or otherwise) on the figure-eight network (see Example 1), then it is also representable on any network as in Figure 3(j). However, the first network is not embedded in the wide sense in the second one and vice versa.

3 Topological Properties

A (two-terminal) network G has the *(equilibrium-) existence property* for weighted network congestion games if every such game on G has at least one pure-strategy (Nash) equilibrium. G has the stronger *finite improvement property* for weighted network congestion games if every such game on it moreover has the finite improvement property. As the following lemma shows, both properties of networks are "hereditary".

Lemma 2. If a network has the existence property, then so does every network homeomorphic to it and every network embedded in the wide sense in it. The same is true for the finite improvement property.

Proof. A logically equivalent proposition is the following. If some weighted network congestion game on a network G does *not* possess a pure-strategy equilibrium or the finite improvement property, then a game with the same quality exists on every network homeomorphic to G and on every network in which G is embedded in the wide sense. The existence of such games follows from Lemma 1. ■

The existence property for weighted network congestion games is preserved also by the operation of connecting networks in series. The reason the connected networks bestow the existence property on the resultant network G is that, as the proof of the following lemma shows, any network congestion game on G is the superposition (see Section 2.1) of such games on them.

Lemma 3. A network made of two or more networks connected in series has the existence property if and only if each of the constituent networks has that property.

Proof. Let G be a network made of m (≥ 2) networks, $G_1, G_2, \dots, G_m$, connected in series. For every weighted network congestion game Γ on G and for each player, choosing an allowable route r in G is equivalent to choosing m allowable routes $r_1, r_2, \dots, r_m$ in $G_1, G_2, \dots, G_m$, respectively, and connecting them in series. Therefore, Γ can be represented as the superposition of m such games – one on each constituent network. In each of the m games, the players and their weights, as well as the cost function and the allowable direction and users for each edge, are as in Γ. This proves that if for $k = 1,2, \dots, m$ every weighted network congestion game on G_k has a pure-strategy equilibrium, this is so also for G.

Conversely, if there is a weighted network congestion game without an equilibrium on G_k, for some $1 \leq k \leq m$, then such a game exists also on G. Specifically, the superposition of the game on G_k and any games on the other $m - 1$ networks (say, games with zero payoffs) is (isomorphic to) a game on G that does not have an equilibrium. ■

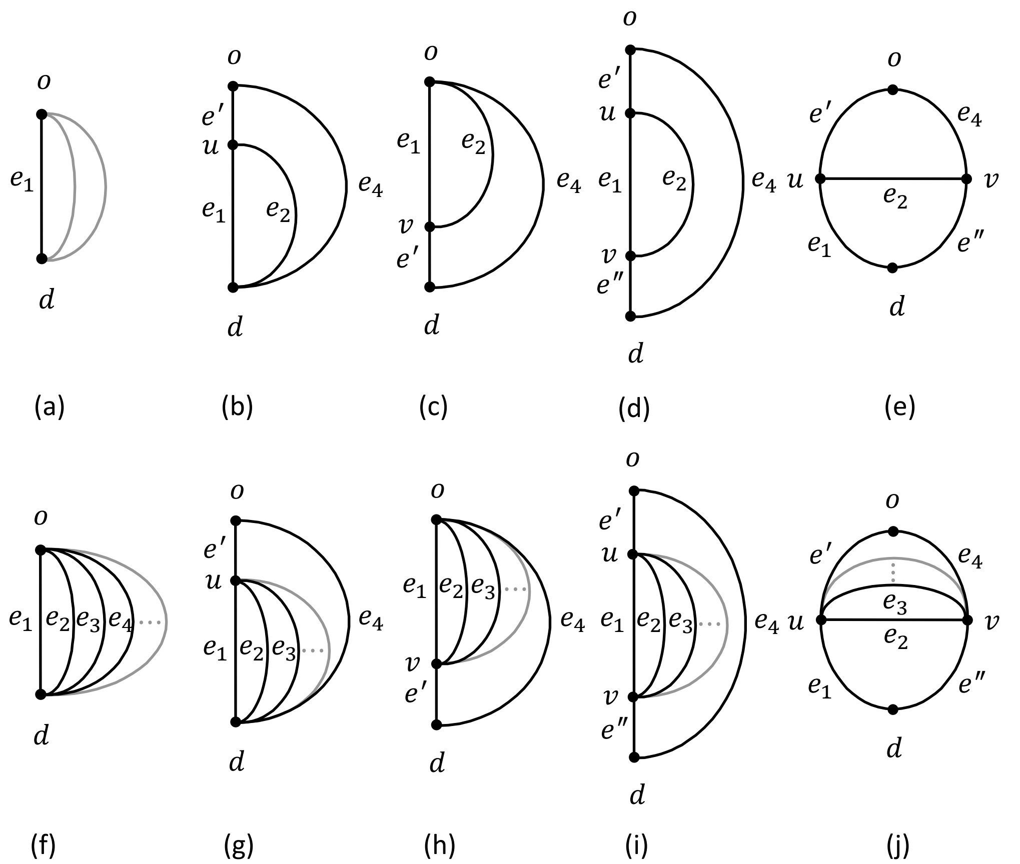

Figure 3. A two-terminal network homeomorphic to any of those depicted here is said to be nearly parallel. A gray, unmarked curve indicates an optional edge and a gray ellipsis mark indicates any number of such edges. The networks in (a)–(f) have the existence property for weighted network congestion games, which means that every such game on them has a pure-strategy equilibrium. The networks in (g)–(j) lack this property.

For the finite improvement property, a result similar to Lemma 3 does not hold. Indeed, as the next theorem shows, virtually the only networks with this property are the *parallel (-link) networks*, which are the networks that have only one edge or are made of several single-edge networks connected in parallel (Figure 3(a) and (f)).

Theorem 2. For a two-terminal network G, the following conditions are equivalent:

(i) Every weighted network congestion game on G has the finite improvement property.

(ii) G is homeomorphic either to a parallel network or to a parallel network connected in series with one or two single-edge networks.

The main result of this paper is the following theorem, which identifies all networks with the existence property for weighted network congestion games. As it shows, essentially the only non-parallel such networks are the networks that can be obtained from the three-edge parallel network by "relocating" one edge's end vertices.

Theorem 3. For a two-terminal network G, the following conditions are equivalent:

(i) Every weighted network congestion game on G has a pure-strategy equilibrium.

(ii) G is homeomorphic to one of the networks in Figure 3(a)–(f) or to a network made of several such networks connected in series.

(iii) None of the networks in Figure 3(g)–(j) or in Figure 4 is embedded in the wide sense in G.

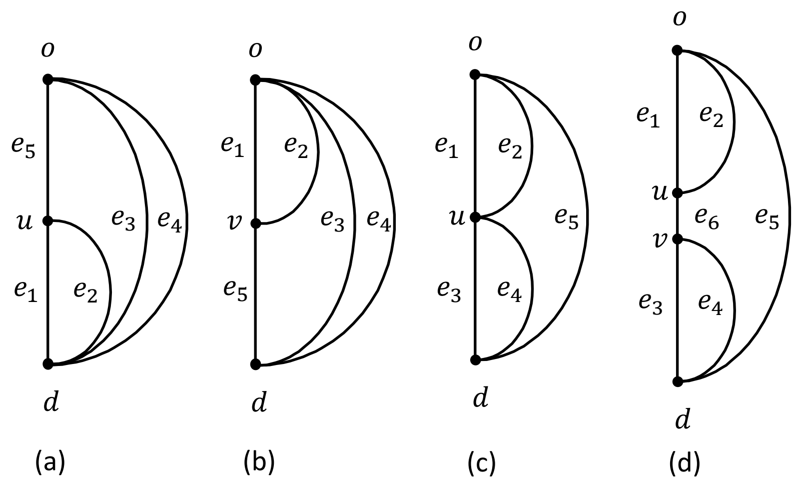

Figure 4. The forbidden networks. On each of these networks, there are weighted network congestion games without pure-strategy equilibria.

Condition (ii) and (iii) in the theorem give two alternative characterizations for the set $\mathcal{G}$ of all networks with the existence property for weighted network congestion games. The difference between them is that (iii) directly characterizes the set $\mathcal{G}^{\mathsf{C}}$ of all networks *without* that property. Its list of networks can be shortened to only six items: the five-edge versions of the networks in Figure 3(g) and (h) (both of which are embedded in the wide sense in each of the networks in (i) and (j)) and the four networks in Figure 4. These networks are the *minimal* elements in $\mathcal{G}^{\mathsf{C}}$, i.e., those in which no other element is embedded in the wide sense. As the existence property is hereditary (Lemma 2), the set of all networks lacking it is completely specified by its minimal elements.

As an illustration of Theorem 3, consider the (underlying undirected) network in Figure 2(a). That network G does not satisfy condition (ii) because (1) unlike the networks in Figure 3(a)–(f), it has four vertices of degree three, and (2) it clearly cannot be constructed by connecting any two networks in series. Therefore, there are weighted network congestion games on G without pure-strategy equilibria. This conclusion also follows from the fact that *every* 2×2 game can be represented as a weighted network congestion game on G (Milchtaich 2013).

As evident from its normal form, the specific game shown in Figure 2(a) does have a pure-strategy equilibrium, indeed, a dominant-strategy one. Interestingly, however, it can be shown that a weighted network congestion game with the same normal form does not exist on any two-terminal network with the existence property. Thus, the existence of a pure-strategy equilibrium in that simple, symmetric 2×2 normal-form game cannot be linked to topological equilibrium-existence. Viewed from a wider perspective, this finding is not surprising. In any game, any (pure) strategy profile can be made an equilibrium by simply boosting the associated payoffs. There is in general no reason to expect this “local” change to bring about representability as a network congestion game on a particular kind of network, which is a “global” property of the game in the sense of depending on all payoffs.

The proofs of Theorems 2 and 3 are given below. They are based on the following graph theoretic result (Milchtaich 2005, Proposition 2.1), which relates all (two-terminal) networks to two special kinds of networks. A *nearly parallel network* is any network that either has only one route or can be constructed by: (1) connecting two single-route networks in parallel, (2) adding any number of edges with identical end vertices and, finally,

(3) subdividing each of these edges any number of times. Depending on whether at most one edge or several edges were added in the second step, a nearly parallel network is homeomorphic to one of those on the upper or lower row, respectively, in Figure 3. The *forbidden networks* are the four specific networks depicted in Figure 4.

Proposition 1. For every two-terminal network G, one, and only one, of the following conditions holds:

(i) G is nearly parallel or it is made of two or more nearly parallel networks connected in series.
(ii) One of the forbidden networks is embedded in the wide sense in G.

3.1 The finite improvement property

One direction of the equivalence in Theorem 2 is essentially the following well-known result (Milchtaich 1996).

Lemma 4. Every weighted network congestion game Γ on a parallel network has the finite improvement property.

Proof. The following argument (Even-Dar et al. 2003; Fabrikant et al. 2004) identifies a specific (generalized ordinal) potential for Γ. Since the number m of all possible payoffs in the game is finite, they can be listed in an ascending order. Associate with each strategy profile an m-tuple, in which the jth entry ($j = 1,2, \dots, m$) is the number of players whose payoff is the jth entry in the list of possible payoffs. Next, rank all strategy profiles lexicographically with respect to these m-tuples. Thus, the highest-ranking strategy profile has the smallest number of players receiving the lowest possible payoff, and in case of a tie, the smallest number of players with the second-lowest payoff among the tied strategy profiles, and so on. The function P that associates each strategy profile with its rank is a potential for Γ. Whenever a single player i unilaterally changes his strategy and increases his payoff as a result, the new strategy profile is ranked higher than the original one. This is because the only players negatively affected by player i's move are those using his new strategy, and their new payoff is equal to i's new payoff and thus higher than his old one. ■

To prove the other direction of the equivalence in Theorem 2, the following two examples are needed.

Example 1. Two players, with weights $w_1 = 1$ and $w_2 = 2$, choose routes in the figure-eight network where edges e_1 and e_2 are connected in parallel, edges e_3 and e_4 are connected in parallel, and the two pairs are connected in series. The cost functions are $c_{e_1}(x) = c_{e_2}(x) = \sqrt{x}$ and $c_{e_3}(x) = c_{e_4}(x) = 0.35x$. Suppose that the two players take turns in changing their strategies, with player 1 alternating between e_1e_3 and e_2e_4 and player 2 alternating between e_1e_4 and e_2e_3. If the order of moves is such that player 2 "chases" player 1, meaning that he moves from the first to the second strategy or back right after player 1 does the same, all changes of strategy are beneficial. Thus, the game does not have the finite improvement property. It does however have four equilibria, which are the strategy profiles where each edge is used by a single player. Note that the example can be immediately extended to an "expanded" figure-eight network, in which the two pairs of parallel edges are separated by an additional edge (as in Figure 4(d) with e_5 removed).

Example 2. Three players, with weights $w_1 = w_2 = 1$ and $w_3 = 2$, choose routes in the network in Figure 3(b) or (c). The cost functions are given (for $1 \leq x \leq 4$) by $c_{e_1}(x) = 8x$, $c_{e_2}(x) = 3x + 6, c_{e_4}(x) = 32.75 - 9/x$ and $c_{e'}(x) = 8x - x^2$. It is not difficult to check that, starting with the strategy profile in which the routes of players 1 and 2 include e_1 and that of player 3 includes e_2, the following is an improvement cycle: player 1 moves to use e_2, player 2 moves to use e_4, player 3 also moves to use e_4, player 1 moves back to e_1, player 2 does the same, and player 3 moves back to use e_2, thus completing the cycle. (The cycle can moreover be made a best-improvement one simply by not allowing each player to use the single edge he does not actually use.) Note that an equilibrium would be immediately reached if player 2, rather than 1, moved first (to e_4), and a different equilibrium would be reached if player 3, rather than 2, moved second (to e_4).

Proof of Theorem 2. If G satisfies (ii), then by Lemmas 2 and 4 it also satisfies (i). If G satisfies (i), then by Lemma 2 and Examples 1 and 2 none of the following networks is embedded in the wide sense in it: the figure-eight network or its "expanded" version (see Example 1), the networks in Figure 3(b) and (c), and hence also all the other non-parallel networks in Figure 3 and Figure 4 (in each of which one of the last two networks is embedded in the wide sense). By Proposition 1, this conclusion implies (ii). ■

3.2 Networks with the existence property

By Lemmas 3 and 4, a network made of several parallel networks connected in series has the existence property for weighted network congestion games. However, none of the networks in Figure 3(b)–(e) is of this kind. Indeed, the one in (e), dubbed the *Wheatstone network*, is not even *series-parallel*, meaning that it cannot be constructed from single-edge networks by *any* sequence of operations of connecting networks in series or in parallel. Thus, establishing the existence property for these networks requires a different approach.

Somewhat unintuitively, the first step in significantly extending topological equilibrium-existence beyond parallel networks is establishing it for a special kind of parallel networks, namely, those in Figure 3(a), but for a larger class of games, where it is possible for a player's weight to only impact the cost for the *other* players. In a weighted network congestion game *in the wide sense* (see example in Figure 2(b)), each edge e is associated with a pair of nondecreasing cost functions, $c_e\colon (0, \infty) \to (-\infty, \infty)$ and $d_e\colon [0, \infty) \to (-\infty, \infty)$, and its cost for each player i is given by

$$c_e(f_e) + d_e(f_e - w_i).$$

The second term differs from the first one in not involving *self-effect*: the argument $f_e - w_i$ is the total weight of the other users of e, excluding player i himself. Lack of self-effect may entail that the cost of an edge is higher for lower-weight players than for higher-weight ones. Thus, as in a network congestion game with player-specific costs, an edge's cost is not necessarily the same for all users. Parenthetically, for weighted network congestion games in the wide sense, a result similar to Lemma 4 does not hold. Indeed, the finite improvement (and even best-improvement) property is not guaranteed even for a three-edge parallel network. Nevertheless, as the following lemma shows, an equilibrium always exists for that network.

Lemma 5. Every weighted network congestion game in the wide sense Γ on a parallel network G with three or fewer edges has a pure-strategy equilibrium.

Proof. Assume, without loss of generality, that G has precisely three edges (some of which may not be allowable to any player), and hence three routes. Identify the edges with three points on an imaginary cycle and say that edge e *follows* edge e' (which *precedes* e) if e is the first edge encountered with when moving along the cycle from e' in the clockwise direction. There are two possible cases: either no player has more than two allowable edges, or at least one player is allowed to use all three. The analysis of both cases uses the following simple result.

Claim 1. Let e and e' be two edges in G that are both allowable to two players i and j with $w_i \leq w_j$. If both players use e and i would not benefit from unilaterally moving to e', then the same is true for j.

This follows from the monotonicity of the cost functions $c_{e'}$ and d_e, which implies that if

$$c_{e'}(f_{e'} + w_i) + d_{e'}(f_{e'}) \geq c_e(f_e) + d_e(f_e - w_i),$$

then a similar inequality holds with i replaced by j.

First case: No player is allowed to use all edges. Associate with each strategy profile (which assigns an edge in G to each player) the total weight $\widehat{w}$ of the players whose edge follows another edge that is allowable for them. There is obviously a unique strategy with $\widehat{w} = 0$, which trivially satisfies the following:

> Each of the players is either not allowed to or would not benefit from moving from his current edge to the preceding edge. (Q)

Claim 2. For every strategy profile satisfying Q that is not an equilibrium, there is a best-improvement path that starts at that strategy profile and ends at another strategy profile satisfying Q with a higher $\widehat{w}$.

To prove Claim 2, consider a strategy profile satisfying Q such that the cost to some player i can be reduced by moving i to some (allowable) edge e, which is necessarily the one following (rather than preceding) his current edge e'. Such a move creates a strategy profile with a higher flow on e and a lower flow on e'. This strategy profile may or may not have property Q. However, due to the monotonicity of the cost functions, Q does not hold only if, for one or more of the players using e, moving to (the preceding) edge e' is both allowed and beneficial. If this is so, choose one of these players with the highest weight, move that player from e to e', and repeat doing so until no more players can benefit from this move. Clearly, player i is not one of the movers. Indeed, his incentive to return to e' can only decrease with each move. Therefore, Claim 1 implies that $w_i > w_j$ for each of the movers j. Thus, the strategy profile reached after the last move differs from the original one in that player i uses e rather than e' and the opposite is true for a certain number (possibly, zero) of other players. The total weight w' of the latter must satisfy $w' < w_i$. Otherwise (that is, if $w' \geq w_i$), for each of them j, the monotonicity of the cost functions and the fact that $w_j < w_i$ would imply the following:

$$(c_{e'}(f_{e'}) + d_{e'}(f_{e'} - w_j)) - (c_e(f_e + w_j) + d_e(f_e))$$
$$\geq (c_{e'}(f_{e'} - w' + w_i) + d_{e'}(f_{e'} - w')) - (c_e(f_e + w') + d_e(f_e + w' - w_i)).$$

However, the left-hand side is (strictly) negative at least for the player j who was the last to move from e to e' (otherwise the move would not have benefited him), while the right-hand side is positive since it gives the reduction in the cost for i when he moved from e' to e. This shows that the above inequality, and hence also $w' \geq w_i$, cannot hold.

The result that $w_i - w'$ is positive means that f_e is higher, and $f_{e'}$ is lower, than the respective flow in the original strategy profile, before i moved. Since the flow on the third edge e'' did not change, it follows that there are still no players who would gain from moving from e'' to (the preceding edge) e or from moving from e' to (the preceding edge) e''. Hence, Q holds for the new as well as for the original strategy profile. In the former, the total weight $\widehat{w}$ of the players whose current edge follows another allowable edge is higher by $w_i - w'$ than in the latter. This completes the proof of Claim 2.

Since $\widehat{w}$ is bounded by the total weight of all players, Claim 2 proves that, if no player is allowed to use more than two edges, an equilibrium exists.

Second case: Γ has some players i with three allowable edges, possibly in addition to players j with only one or two such edges. Re-index the players in the game in such a way that, for some $1 \leq k \leq n$, the inequalities $i \leq k < j$ hold for all players i and j as above (who differ in their number of strategies) and $w_i \leq w_j$ holds for *all* i and j with $i < j$. (The cardinality assumption in the definition of weighted network congestion game implies that such re-indexing is possible.) For each player i, define Γ_i as the game obtained from Γ by excluding i and all lower-index players, so that they do not choose routes and do not contribute to the flows. In addition, define $\Gamma_0 = \Gamma$.

It follows from the first part of the proof that Γ_k, whose set of players (which may be empty) consists of all the players in Γ with one or two allowable edges, has a pure-strategy equilibrium. To prove that such an equilibrium exists also in Γ it suffices to show that, for all $1 \leq i \leq k$, the existence of an equilibrium in Γ_i implies the same for Γ_{i-1}. In fact, for any equilibrium in Γ_i, simply choosing a best response strategy for player i gives an equilibrium in Γ_{i-1}. Clearly, any player j whose edge is different from the edge e chosen by i still cannot gain from changing his strategy. (His incentive to do so is, if anything, even lower than before.) The same is true if j's strategy is e. Since $w_i \leq w_j$, and since moving from e to any other edge e' is not beneficial to i, by Claim 1 the same applies to j. ∎

The significance of Lemma 5 lies in the fact that every weighted network congestion game in the wide sense Γ (and, in particular, every such game in the "regular" sense) on any of the non-parallel networks in Figure 3(b)–(e) is similar (see Section 2.1) to such a game on a parallel network as in (a). That game is obtained from Γ by a procedure dubbed *parallelization*, which is described in the proof of the following lemma. Parallelization both changes the network to a parallel one and transforms some cost functions with self-effect (c_e's) into cost functions without self-effect (d_e's) and vice versa. This suggests that the two kinds of cost functions may be intimately connected.

Lemma 6. Every weighted network congestion game in the wide sense Γ on any of the networks G in Figure 3(b)–(e) is similar to such a game $\tilde{\Gamma}$ on a parallel network with three edges.

Proof. Let $\tilde{G}$ be the parallel network, with edges e_1, e_2 and e_4, that is obtained from G by *contracting* edge e' and, if G has a fifth edge e'', also contracting that edge. Contraction (Diestel 2005) is the one-sided inverse of the operation of terminal subdivision (Figure 1(c)): it eliminates the edge and its non-terminal end vertex. Each of the three routes in $\tilde{G}$ corresponds to a route in G, which includes the former's single edge and traverses it in the same direction. This correspondence between routes is one-to-one and onto, with one exception. The single exception is route $e_4e_2e_1$ in the Wheatstone network (Figure 3(e)), which does not have a corresponding route in the parallel network $\tilde{G}$. However, that route may be ignored since, by symmetry, it suffices to consider network congestion games on the Wheatstone network in which the allowable direction of e_2 is from u to v. Thus, it suffices to consider games Γ on G in which every route that is allowable for some player has a corresponding (single-edge) route in $\tilde{G}$. The next step is to describe the corresponding game $\tilde{\Gamma}$ on $\tilde{G}$.

The following description concerns the case in which G is the Wheatstone network. The other three cases (Figure 3(b)–(d)) are rather similar (indeed, somewhat simpler). The game $\tilde{\Gamma}$ on $\tilde{G}$ inherits from Γ its set of players, their weights and their strategy sets (with the identification of routes in G and $\tilde{G}$ described above). The cost functions in $\tilde{\Gamma}$ (which in the following are marked by a tilde accent) are derived from those in Γ (unaccented) as follows. For $0 \leq y < x \leq w$, where $w = \sum_i w_i$ is the players' total weight,

$$\begin{aligned} \tilde{c}_{e_1}(x) &= c_{e_1}(x) - d_{e''}(w-x), & \tilde{d}_{e_1}(y) &= d_{e_1}(y) - c_{e''}(w-y), \\ \tilde{c}_{e_2}(x) &= c_{e_2}(x), & \tilde{d}_{e_2}(y) &= d_{e_2}(y), \\ \tilde{c}_{e_4}(x) &= c_{e_4}(x) - d_{e'}(w-x), & \tilde{d}_{e_4}(y) &= d_{e_4}(y) - c_{e'}(w-y). \end{aligned}$$

It has to be shown that the games Γ and $\tilde{\Gamma}$ are similar. That is, for each player i, the difference between the costs to i in Γ and in $\tilde{\Gamma}$ can be expressed as a function of the route choices of the other players. If i's route includes e_2 (hence, does not include e_1 or e_4), the difference can be written as

$$c_{e'}(w_{-i,-4} + w_i) + d_{e'}(w_{-i,-4}) + c_{e''}(w_{-i,-1} + w_i) + d_{e''}(w_{-i,-1}), \tag{2}$$

where $w_{-i,-4}$ or $w_{-i,-1}$ is the total weight of the players other than i whose route does not include e_4 or e_1, respectively. The same expression gives the difference between the costs in Γ and in $\tilde{\Gamma}$ also if i's route does include either e_1 or e_4. For example, if the route includes e_1, its total cost for i in Γ is $c_{e'}(w - f_{e_4}) + d_{e'}(w - f_{e_4} - w_i) + c_{e_1}(f_{e_1}) + d_{e_1}(f_{e_1} - w_i)$, and in $\tilde{\Gamma}$ the cost is $\tilde{c}_{e_1}(f_{e_1}) + \tilde{d}_{e_1}(f_{e_1} - w_i)$. It is not difficult to see that the difference between the two costs can again be written as (2). Thus, the difference is independent of i's route, as had to be shown. ∎

Parenthetically, the assertion of Lemma 6 cannot be strengthened to isomorphism between Γ and $\tilde{\Gamma}$. In other words, the collection of all finite games representable as weighted network congestion games in the wide sense on the networks in Figure 3(b)–(e) is a proper superset of those representable using (a). For example, it is not difficult to show that the 2×2 game in Figure 2(b) cannot be represented as a weighted network congestion game in the wide sense on any parallel network; no such game shares its normal form.

An immediate corollary of the last two lemmas is the following result, which together with Lemmas 2, 3 and 4 establishes the "positive" part of Theorem 3.

Lemma 7. Every weighted network congestion game in the wide sense on one of the networks in Figure 3(a)–(e) has a pure-strategy equilibrium.

3.3 Networks without the existence property

A network without the existence property for weighted network congestion games can be obtained from any network homeomorphic to one of those in Figure 3(b)–(e) by simply adding *any* single edge. This is because a network obtained this way necessarily has one (or more) of those in Figure 3(g)–(j) or Figure 4 embedded in the wide sense in it. As the following five examples show, there are four-player weighted network congestion games on the networks in Figure 3(g)–(j) and three-player games on those in Figure 4 that do not have pure-strategy equilibria. It can moreover be shown that, with one possible exception, these numbers of players are minimal for non-existence of equilibrium. Specifically, every three-player weighted network congestion game in the wide sense on any of the networks in Figure 3(g)–(i) has a pure-strategy equilibrium, and the same is true for every two-player such game on any of the networks in Figure 4.

Example 3. Four players, with weights $w_1 = 1$, $w_2 = 2$ and $w_3 = w_4 = 3$, choose routes in one of the networks in Figure 3(g)–(j). Each player has two allowable routes, each of which includes exactly one of the edges e_1, e_2, e_3 and e_4. One route for player 1, 2, 3 and 4, referred to as Left, includes e_2, e_2, e_1 and e_3, respectively, and the other route, Right, includes e_3, e_4, e_2 and e_4, respectively. The cost functions of edges e_1, e_2, e_3, e_4 and e' are positive and satisfy: $c_{e_1}(3) = 16;\ c_{e_2}(1) = 2, c_{e_2}(3) = 3, c_{e_2}(4) = 15, c_{e_2}(5) = 17;\ c_{e_3}(1) = 4, c_{e_3}(3) = 10, c_{e_3}(4) = 14;\ c_{e_4}(2) = 9, c_{e_4}(3) = 18, c_{e_4}(5) = 19;\ c_{e'}(6) = 1, c_{e'}(7) = 6, c_{e'}(9) = 7$. The cost of edge e'', if it exists in the network, is 0. It can be verified that Left is the better choice for player 3, player 1 or player 4 if and only if the strategy of player 2, player 3 or player 1, respectively, is also Left. Therefore, in any equilibrium where player 2 plays Left or Right, the other players necessarily do the same. However, this means that in the first case player 2 can decrease his cost from 10 to 9 by (unilaterally) changing his choice to Right, and in the second case, he can decrease it from 19 to 18 by changing to Left. This contradiction proves that a pure-strategy equilibrium does not exist.

In all but one of the networks considered in Example 3, there is only one legitimate way to direct the edges. The exception is the network in Figure 3(j), where e_2 and e_3 may have identical or opposite directions. In Example 3, the former holds, and in the next example, the latter holds. This proves that, for Figure 3(j), and trivially also for all the other networks in Figure 3 and Figure 4, directionality is not an important consideration: pre-determining the edges' directions would not affect equilibrium existence.

Example 4. This example differs from the previous one in that it only refers to Figure 3(j) and in that the players' routes are different: Left for player 1, 2, 3 and 4 means $e'e_2e''$, $e'e_1$,

$e_4e_3e_1$ and $e'e_1$, respectively, and Right means e_4e'', e_4e'', $e'e_1$ and $e'e_2e''$, respectively. (Note that for player 3 the two routes are actually on the sides opposite to those suggested by their names.) In addition, the cost functions are different, and satisfy: $c_{e_1}(5) = 5$, $c_{e_1}(6) = 16, c_{e_1}(8) = 17;\ c_{e_2}(1) = 1, c_{e_2}(3) = 2, c_{e_2}(4) = 6;\ c_{e_3}(3) = 1;\ c_{e_4}(3) = 5$, $c_{e_4}(4) = 10, c_{e_4}(5) = 15;\ c_{e'}(6) = 3, c_{e'}(7) = 6, c_{e'}(8) = 12;\ c_{e''}(3) = 4, c_{e''}(4) = 12, c_{e''}(6) = 13$. It can be verified that Left is the better choice for player 3, player 1 or player 4 if and only if the strategy of player 2, player 3 or player 1, respectively, is also Left. Therefore, in any equilibrium where player 2 plays Left or Right, the other players necessarily do the same. However, this means that in the first case player 2 can decrease his cost from 20 to 19 by (unilaterally) changing his choice to Right, and in the second case, he can decrease it from 18 to 17 by changing to Left. This contradiction proves that a pure-strategy equilibrium does not exist.

Example 5. Three players, with weights $w_1 = 3$ and $w_2 = w_3 = 4$, choose routes in the network in Figure 4(a) or (b). The only restrictions on route choices are that edge e_2 is only allowable to player 2, who is not allowed to use e_1, and e_4 is only allowable to player 3, who is not allowed to use e_3. Thus, there are two allowable routes for each player: Left, which includes e_5, and Right, which does not. The costs of the two private edges satisfy $c_{e_2}(4) = 7$ and $c_{e_4}(4) = 13$. Those of the other edges are given (for $x \geq 3$) by $c_{e_1}(x) = x$, $c_{e_3}(x) = 15 - 0.75(2 - 0.25x)^9$ and $c_{e_5}(x) = x$. It can be verified that Left is the better choice for player 1, player 2 or player 3 if and only if the strategy of player 2, player 3 or player 1, respectively, is Right. It follows that a pure-strategy equilibrium does not exist.

Example 6. Three players, with weights $w_1 = 1$ and $w_2 = w_3 = 2$, choose routes in the network in Figure 4(c). The only restrictions are that edge e_2 is only allowable to player 2, who is not allowed to use e_1, and e_4 is only allowable to player 3, who is not allowed to use e_3. Thus, there are two allowable routes to each player: Left, which does not include e_5, and Right, which does. The costs of the two private edges satisfy $c_{e_2}(2) = 3$ and $c_{e_4}(2) = 9$, and those of the other edges satisfy: $c_{e_1}(1) = 1, c_{e_1}(2) = 2, c_{e_1}(3) = 8;\ c_{e_3}(1) = 2, c_{e_3}(2) = 10, c_{e_3}(3) = 12;\ c_{e_5}(x) = 4x$. It can be verified that Left is the better choice for player 1, player 2 or player 3 if and only if the strategy of player 2, player 3 or player 1, respectively, is Right. It follows that a pure-strategy equilibrium does not exist.

Example 7. Three players, with weights $w_1 = 1$, $w_2 = 5$ and $w_3 = 10$, choose routes in the network in Figure 4(d). The only restrictions are that edge e_2 is only allowable to player 2, who is not allowed to use e_1, and e_4 is only allowable to player 3, who is not allowed to use e_3. Thus, there are two allowable routes for each player: Left, which does not include e_5, and Right, which does. Three of the edges have constant costs, $c_{e_2} = 1.3$, $c_{e_4} = 6.25$ and $c_{e_5} = 40$, and three have increasing costs, $c_{e_1}(x) = 2x, c_{e_3}(x) = 5x$ and $c_{e_6}(x) = 3.55\sqrt{x}$. It can be verified that Left is the better choice for player 1, player 2 or player 3 if and only if the strategy of player 2, player 3 or player 1, respectively, is Right. It follows that a pure-strategy equilibrium does not exist.

Another example of a game without a pure-strategy equilibrium on the network in Figure 4(d) can be obtained from Example 6 by simply setting $c_{e_6} = 0$.

Theorem 3 can now be proved.

Proof of Theorem 3. If condition (i) holds, then it follows from Lemma 2 and Examples 3, 5, 6 and 7 that condition (iii) also holds. If condition (ii) holds, then it follows from Lemmas 2, 3, 4 and 7 that condition (i) holds. It remains to observe that, by Proposition 1, condition (iii) implies (ii). ■

For weighted network congestion games in the wide sense, a result very similar to Theorem 3 holds, except that the networks in Figure 3(f) are removed from condition (ii) and added to the list in (iii). The proof is very similar to that above, but also uses the following example, which is obtained from Example 3 by parallelization (see the proof of Lemma 6).

Example 8. Four players, with weights $w_1 = 1$, $w_2 = 2$ and $w_3 = w_4 = 3$, choose routes in the network in Figure 3(f). Each player has two allowable routes: Left, which for player 1, 2, 3 and 4 means e_2, e_2, e_1 and e_3, respectively, and Right, which means e_3, e_4, e_2 and e_4, respectively. The costs of the edges satisfy: $c_{e_1}(3) = 16;\ c_{e_2}(1) = 2, c_{e_2}(3) = 3, c_{e_2}(4) = 15, c_{e_2}(5) = 17;\ c_{e_3}(1) = 4, c_{e_3}(3) = 10, c_{e_3}(4) = 14;\ c_{e_4}(2) = 2, c_{e_4}(3) = 11, c_{e_4}(5) = 12$. In addition, $d_e = 0$ for all edges e except e_4, for which $d_{e_4}(0) = 0, d_{e_4}(2) = 1$, $d_{e_4}(3) = 6$. It can be verified that Left is the better choice for player 3, player 1 or player 4 if and only if the strategy of player 2, player 3 or player 1, respectively, is also Left. Therefore, in any equilibrium where player 2 plays Left or Right, the other players necessarily do the same. However, this means that in the first case player 2 can decrease his cost from 3 to 2 by (unilaterally) changing his choice to Right, and in the second case, he can decrease it from 18 to 17 by changing to Left. This contradiction proves that a pure-strategy equilibrium does not exist.

4 Related Models and Open Problems

Existence of pure-strategy equilibrium and the finite improvement property are two of several properties of network congestion games that can be linked to the network topology. A third one is the property that all pure-strategy equilibria in the game are *strong*. Holzman and Law-yone (1997, 2003) studied this property in the context of unweighted network congestion games in which all edges are public. They showed that all such games on a network have the above property if and only if the network is *extension-parallel*, meaning that it can be built from single-edge networks by repeatedly connecting networks in series or in parallel, with the proviso that in the first case at most one network can have more than one edge. An equivalent way of stating this result is that an extension-parallel network is a necessary and sufficient condition for *weak Pareto efficiency* of all equilibria in all corresponding games, meaning that it is never possible to alter the players' equilibrium route choices in a way that benefits them all. The equivalence holds because an equilibrium is strong if and only if the strategy choices of every subset of players constitute a weak Pareto efficient equilibrium in the subgame defined by fixing the strategies of the remaining players. That subgame is itself an unweighted network congestion game with public edges on the same network.

Holzman and Law-yone's result was originally established for *directed* networks, that is, with the edges' directions fixed as part of the network's specification. However, it holds also in the present setting of undirected networks, where the edges' directions may vary with the game considered. An undirected network is extension-parallel if and only if it has *linearly independent* routes, in the sense that each route includes at least one edge that is not part of any other route (Milchtaich 2006a).

A similar connection between the network topology and the weak Pareto efficiency of all equilibria holds for *nonatomic* network congestion games with a continuum of identical players (Milchtaich 2006a). That is, a necessary and sufficient condition for weak Pareto efficiency of all equilibria in all such games on an (undirected) network (regardless of the cost functions and the directions that the game assigns to the edges) is that the network has linearly independent routes. Moreover, unlike in the finite case, this result holds also with non-identical players, that is, with player-specific cost functions.

A network has the *uniqueness property* for a particular variety of network congestion games if in every game of that kind on the network the players' (pure-strategy) equilibrium costs are unique. This topological property is not relevant for finite network congestion games, where it is virtually impossible to guarantee uniqueness, or for nonatomic ones with identical cost functions, where the equilibrium costs are always unique. For nonatomic network congestion games with player-specific costs, a network has the uniqueness property if and only if it is nearly parallel (Figure 3) or consists of two or more nearly parallel networks connected in series (Milchtaich 2005). The complementary class of networks that allow for multiple equilibrium costs consists of all networks in which one of the forbidden networks (Figure 4) is embedded in the wide sense. A similar result holds for network congestion games with finitely many players in which flow is splittable among multiple routes (Richman and Shimkin 2007).

The topological efficiency and uniqueness problems for nonatomic network congestion games are not directly related to the topological equilibrium-existence problem studied in this paper, which concerns finite games. (For nonatomic network congestion games, the existence of pure-strategy equilibrium is not an issue, since it is guaranteed by weak assumptions on the cost functions; see Schmeidler 1970.) Nevertheless, the solutions to the three problems turn out to have broadly similar forms. In particular, each topological property is equivalent to the nonexistence of an embedded (in the wide sense) network belonging to a particular short list of "bad" networks. The solutions are also all formulated in terms of undirected networks, which may attest to the practical merit of viewing directionality as belonging to the game rather than pre-determined by the network. However, this perspective leaves open the following question: For which directed networks is the existence of equilibrium guaranteed in all weighted network congestion games that respect the edges' directions? The remarks that precede Example 4 may be the first step in answering this question.

The rest of this section considers several other models that are related to but different from that studied in Section 3, and presents several results and open problems pertaining to these models.

4.1 Public edges

The existence property for weighted network congestion games where all players are allowed to use all edges is less demanding than in the general case considered above (Theorem 3). In particular, it holds for the nearly parallel networks in Figure 3(g)–(i). This result is proved in Milchtaich (2006b) by first showing that on a parallel network (even one with four edges or more; Figure 3(f)) every weighted network congestion game in the wide sense with public edges has a pure-strategy equilibrium. Indeed, an equilibrium can easily be found by employing the greedy best response algorithm (Fotakis et al. 2006), whereby the players enter the game one by one with heavier players entering first (see the proof of Lemma 5, second case). A straightforward generalization of the parallelization argument used in Lemma 6 then extends the result to all the networks in Figure 3(a)–(i). (The argument partially applies also to the remaining nearly parallel networks, which are represented by Figure 3(j). However, it only applies to games in which the edges with end vertices u and v all have the same allowable direction: from u to v or vice versa.)

The main open problem regarding the topological equilibrium-existence problem for weighted network congestion games with public edges is whether, or to what extent, the existence property holds for networks that are not nearly parallel or made of several such networks connected in series. In particular, it is not known whether any of the forbidden networks has this property. An example of a network (with linearly independent routes) that does not have the existence property can be obtained from the forbidden network in Figure 4(a) by subdividing e_1 and joining the resultant new vertex with d by a new edge. A weighted network congestion game with public edges on that network that does not have a pure-strategy equilibrium is presented in Milchtaich (2006b). Another network without the existence property is shown in Figure 5.

4.2 Player-specific costs

Another open problem is the characterization of the networks with the existence property for (unweighted) network congestion games with player-specific costs. It is known that these include all parallel networks (Milchtaich 1996). Although a game of this kind on a parallel network does not always have the finite improvement property, there is a simple algorithm that, starting with any strategy profile, identifies a best-improvement path ending at an equilibrium, whose length is polynomial in the number of players and edges. As in the case of weighted network congestion games with public edges, a parallelization argument extends the equilibrium-existence result to all the nearly parallel networks in Figure 3(a)–(i) (and partially also to (j)) (Milchtaich 2006b).

The set of networks that are known *not* to have the existence property only partially overlaps the corresponding set for weighted network congestion games with public edges. It includes the networks obtained by adding: (1) an edge with end vertices o and u to the network in Figure 4(a) (equivalently, end vertices v and d in (b) or o and d in (c)), (2) an edge with end vertices u and v to the network in Figure 4(d), or (3) an edge with end vertices o and d to the Wheatstone network in Figure 3(e) (Milchtaich 2006b). Moreover, there are network congestion games with player-specific *linear* cost functions (with positive coefficients) on the three resultant networks that do not have pure-strategy equilibria. For the network defined in (3), one example of a (nonlinear) game without an equilibrium is given in Figure 5.

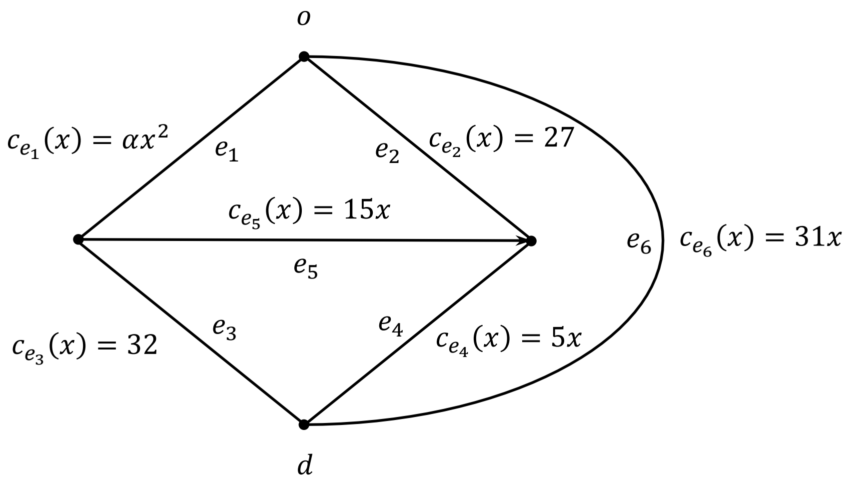

Figure 5. A two-player weighted network congestion game with public edges and weights $w_1 = 1$ and $w_2 = 2$. If the coefficient α in c_{e_1} is $9/8$, there is no pure-strategy equilibrium. However, if $\alpha = 10/9$, there is a unique such equilibrium, in which player 1 chooses the route $e_1 e_5 e_4$ and player 2 chooses $e_1 e_3$ (and is indifferent between that route and $e_2 e_4$). The game can be turned into an (unweighted) network congestion game with player-specific costs by replacing each cost function c_e with a pair of player-specific cost functions such that $c_{1e}(1) = c_e(1)$, $c_{2e}(1) = c_e(2)$ and $c_{1e}(2) = c_{2e}(2) = c_e(3)$. The players' strategies and payoffs in the resultant game are identical to those in the original, weighted one.

4.3 The complexity of the equilibrium-existence decision problem

The desirability of solving the topological equilibrium-existence problem for the kinds of network congestion games considered in the last two subsections is underlined by the fact that, even for a network with only a moderately large number of edges, deciding whether a *specific*, given game has a pure-strategy equilibrium may be computationally difficult. Moreover, as the following proposition shows, this is so even with only two players. The proposition unifies an earlier result of Ackermann and Skopalik (2007), which concerns network congestion games with player-specific costs, and a somewhat stronger version of a result of Dunkel and Schulz (2008), which concerns weighted network congestion games with public edges. The idea of the proof is to start with a simple network congestion game where the unique equilibrium is not strict, and would be eliminated by any small increase in the cost of some edge. That edge is then connected in series with an auxiliary network for which it is difficult to decide whether collision-free routing is possible.

Proposition 2. The problem of deciding whether a pure-strategy equilibrium exists is NP-complete for each of the following two classes of games:

(i) Two-player weighted network congestion games with public edges and nonnegative cost functions.

(ii) Two-player (unweighted) network congestion games with player-specific costs and nonnegative cost functions.

Proof. The hardness of the problem is established by reduction from the directed edge-disjoint paths problem with two pairs of terminal vertices, which is NP-complete (Fortune et al. 1980). The input of that problem is a directed version of a network G similar to that defined in Section 2.2, except that it has two origin vertices, o_1 and o_2, and two destination vertices, d_1 and d_2. It may be assumed that the four terminal vertices are distinct, that there is at least one path beginning with o_1 and ending with d_1 which traverses each of its edges in

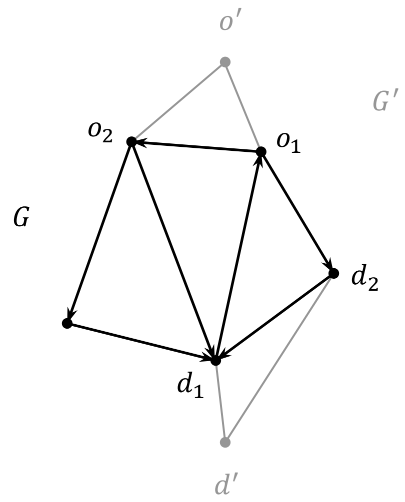

Figure 6. An instance G of the directed edge-disjoint paths problem (black) and the two-terminal network G' obtained from it by the addition of origin and destination vertices, o' and d', and four connecting edges (gray).

the specified direction, a similar path connects o_2 and d_2, and every edge belongs to some path of the first or the second kind. The problem is to decide whether there exists a pair of paths, one connecting o_1 and d_1 and the other connecting o_2 and d_2, that do not share any edges. Turning this decision problem into an equilibrium-existence problem only requires adding to G a few edges and vertices and defining suitable cost functions. The first step is to add new origin and destination vertices, o' and d', and connect each of them with each of the two corresponding original terminal vertices (see Figure 6). The resultant two-terminal network G' is then inserted "between" o and e_1 in the network depicted in Figure 5, so that o is identified with o' and is replaced by d' as an end vertex of e_1. The cost functions in Figure 5 remain in effect, with $\alpha = 10/9$, both (i) in the case in which the two players differ in their weights, $w_1 = 1$ and $w_2 = 2$, and (ii) in the case in which they differ in their cost functions. In case (i), the edges connecting o' with o_j and d' with d_j ($j = 1{,}2$) are assigned the cost function $[x - j]^+$ (where $[y]^+$ means $\max\{y, 0\}$), and in case (ii), the corresponding cost function for player i ($= 1{,}2$) is $[x - j + i - 1]^+$. The remaining edges e, which are those belonging to G, are assigned identical cost functions, which are $c_e(x) = [x - 2]^+/(N + 1)$ in case (i) and $c_{1e}(x) = c_{2e}(x) = [x - 1]^+/(N + 1)$ in case (ii), where N is the number of edges in G.

Each of the cases (i) and (ii) defines a network congestion game belonging to the corresponding class in the proposition. Choosing a route in this game means choosing a route in the network in Figure 5, and if that route includes e_1, also choosing a route in G'. It is easy to see that, for player 2, the cost of a route in G' is 0 if and only if (1) its first and last edges are incident with vertices o_2 and d_2, respectively, and (2) none of the edges in the route is used also by player 1. When (1) holds for player 2's route in G', the cost for player 1 of any route there that begins and ends with the edges incident with o_1 and d_1, respectively, is less than 1, and it is equal to 0 if and only if the route does not share with player 2's route any edges in G. Using these facts, it is not difficult to check that a pair of strategies in the game is an equilibrium if and only if the players' routes in the network in Figure 5 are the indicated equilibrium ones and the costs of their routes in G' are zero. The second condition holds for some pair of routes in G' if and only if the answer to the decision problem specified by G is affirmative.

Note that deciding whether a *given* strategy profile is an equilibrium means checking if each player's route is a least-cost one with respect to the costs determined by the other player's route. The number of required steps is at most of the order of the number of vertices squared. ■

The decision problem considered in Proposition 2 is of course NP-hard also without the restrictive assumptions of public edges and nonnegative costs or with an unbounded number of players. Its hardness is rooted in the fact that, for general networks, the number of routes (and of strategies) may increase exponentially as the number of edges increases. With an unbounded number of players, deciding whether a pure-strategy equilibrium exists may be difficult also with network topologies where the number of routes is comparable with the number of edges.

For network congestion games in which the players may differ in both their weights *and* cost functions, Dunkel and Schulz (2008) showed that the equilibrium-existence decision problem is NP-complete even with parallel networks. The corresponding topological equilibrium-existence problem, by contrast, is quite trivial. On a two-edge parallel network, every weighted network congestion game with player-specific costs has a pure-strategy equilibrium, but this is not so for a three-edge parallel network even in the case of only three players (Milchtaich 1996).

4.4 Matroid congestion games

Each (two-terminal) network topology entails a particular set of combinatorial restrictions on the players' strategy sets in all corresponding network congestion games. For example, for *any* topology, different strategies are *incomparable* in that the set of edges in one strategy is not a subset of that in any other strategy. The restrictions take an extreme form in the case of parallel networks, which correspond to the so-called *singleton congestion games*: each player simply has to choose one allowable edge. This observation leads to the question of whether the existence of equilibrium in the latter and similar classes of network congestion games can be linked directly to the combinatorial structure of the strategy sets, rather than to the network topology giving rise to that structure. Specifically, Ackermann et al. (2009) presented the following combinatorial version of the equilibrium-existence problem: What is the most general combinatorial structure for which a pure-strategy equilibrium is guaranteed to exist in every corresponding congestion game in which players may differ in their weights, and what is that structure when players differ in their cost functions? The congestion games that the two versions of the problem refer to are more general than the corresponding network congestion games considered in this paper. Each player's strategy set is an arbitrary collection of subsets of a common set of resources, which may or may not be the edges of a network.

As Ackermann et al. (2009) showed, the most general games of both kinds for which the existence of equilibrium is guaranteed are *matroid congestion games*, in which the strategy set of each player consists of the bases of a matroid on the set of resources. These games share with singleton congestion games the property (which reflects the corresponding property of bases of a matroid) that all strategies of a player include the same number of resources, but they allow for much more varied and elaborate combinatorial structures, for example, strategy sets that consist of all pairs of resources. However, a noteworthy aspect of these results is that they do not take into account how the strategy sets of different players interweave. This means, in particular, that the existence of a pure-strategy equilibrium in weighted network congestion games and network congestion games with player-specific costs may be guaranteed even if the players share a common strategy set that does *not* consist of the bases of a matroid, for example, if some allowable routes includes fewer edges than others (which is normally the case for the networks in Figure 3(b)–(e)). The results only entail that, with such a strategy set, it is possible to systematically substitute a different edge for each allowable edge for each player, such that with the *modified* strategy sets a pure-strategy equilibrium may not exist. However, a strategy modified in this way is not necessarily a route in the network.

The positive part of the solution to the combinatorial equilibrium-existence problem obtained by Ackermann et al. (2009) does apply to network congestion games. However, its usefulness for the graph-theoretic version studied in the present paper is limited. This assessment is based on the following fact.

Proposition 3. In a network congestion game on a two-terminal network G, the strategy set of a player consists of the bases of a matroid on the set of edges if and only if the sub-network of G that includes only the edges belonging to the player's allowable routes is parallel or is made of several parallel networks connected in series.

Proof. It has to be shown that the first condition (the matroid property) is equivalent to the following graph theoretic one: the player's allowable routes all have the exact same vertices and pass them in the same order. Since different routes have incomparable sets of edges, the routes' sets of edges are the bases of a matroid if and only if they satisfy the bijective exchange axiom (White 1986): there is a one-to-one correspondence between the sets of edges in any pair of allowable routes, such that replacing any edge e in one route with the corresponding edge e' in the other route again gives the set of edges in some allowable route. Clearly, the corresponding edges e and e' must have the same end vertices. Therefore, the bijective exchange axiom is equivalent to the above graph theoretic condition. ■

Acknowledgments

I thank a referee and an Associate Editor for their comments. This research was supported by the Israel Science Foundation (grants No. 1082/06 and 1167/12).

References

Ackermann H, Skopalik A (2007) On the complexity of pure Nash equilibria in player-specific network congestion games. Lect Notes Comput Sci 4858:419–430

Ackermann H, Röglin H, Vöcking B (2009) Pure Nash equilibria in player-specific and weighted congestion games. Theor Comput Sci 410:1552–1563

Diestel R (2005) Graph Theory, 3rd edn. Graduate Texts in Mathematics, Vol. 173. Springer-Verlag, New York

Dunkel J, Schulz AS (2008) On the complexity of pure-strategy Nash equilibria in congestion and local-effect games. Math Oper Res 33:851–868

Even-Dar E, Kesselman A, Mansour Y (2003) Convergence time to Nash equilibria. Lect Notes Comput Sci 2719:502–513

Fabrikant A, Papadimitriou C, Talwar K (2004) The complexity of pure Nash equilibria. In: László Babai (ed) Proceedings of the 36th Annual ACM Symposium on Theory of Computing. pp 604–612

Fischer F, Holzer M, Katzenbeisser S (2006) The influence of neighbourhood and choice on the complexity of finding pure Nash equilibria. Inform Process Lett 99:239–245

Fortune S, Hopcroft J, Wyllie J (1980) The directed subgraph homeomorphism problem. Theor Comput Sci 10:111–121

Fotakis D, Kontogiannis S, Spirakis P (2005) Selfish unsplittable flows. Theor Comput Sci 348:226–239

Fotakis D, Kontogiannis S, Spirakis P (2006) Symmetry in network congestion games: pure equilibria and anarchy cost. Lect Notes Comput Sci 3879:161–175

Goemans M, Mirrokni V, Vetta A (2005) Sink equilibria and convergence. In: Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science. pp 142–154

Holzman R, Law-Yone N (1997) Strong equilibrium in congestion games. Game Econ Behav 21:85–101

Holzman R, Law-yone (Lev-tov) N (2003) Network structure and strong equilibrium in route selection games. Math Soc Sci 46:193–205

Libman L, Orda A (2001) Atomic resource sharing in noncooperative networks. Telecommun Syst 17:385–409

Milchtaich I (1996) Congestion games with player-specific payoff functions. Game Econ Behav 13:111–124

Milchtaich I (2005) Topological conditions for uniqueness of equilibrium in networks. Math Oper Res 30:225–244

Milchtaich I (2006a) Network topology and the efficiency of equilibrium. Game Econ Behav 57:321–346

Milchtaich I (2006b) The equilibrium existence problem in finite network congestion games. Lect Notes Comput Sci 4286:87–98

Milchtaich I (2013) Representation of finite games as network congestion games. Int J Game Theory 42: 1085–1096

Monderer D (2007) Multipotential games. In: Sangal R, Mehta H, Bagga RK (eds) Proceedings of the 20th International Joint Conference on Artificial intelligence. Morgan Kaufmann Publishers, San Francisco, pp 1422–1427

Monderer D, Shapley LS (1996) Potential games. Game Econ Behav 14:124–143

Morris S, Ui T (2004). Best response equivalence. Game Econ Behav 49:260–287

Richman O, Shimkin N (2007). Topological uniqueness of the Nash equilibrium for atomic selfish routing. Math Oper Res 32:215–232

Schmeidler D (1970) Equilibrium points of nonatomic games. J Stat Phys 7:295–300

von Neumann J, Morgenstern O (1953) Theory of Games and Economic Behavior, 3rd edn. Princeton University Press, Princeton, NJ

White N (ed) (1986) Theory of Matroids. Encyclopedia of Mathematics and its Applications. Cambridge University Press, Cambridge, UK