

Network Topology and Equilibrium Existence in Weighted Network Congestion Games

Igal Milchtaich, *Bar-Ilan University*

May 2010

Abstract. Every finite noncooperative game can be presented as a weighted network congestion game, and also as a network congestion game with player-specific costs. In the first presentation, different players may contribute differently to congestion, and in the second, they are differently (negatively) affected by it. This paper shows that the topology of the underlying (undirected two-terminal) network provides information about the existence of pure-strategy Nash equilibrium in the game. For some networks, but not for others, every corresponding game has at least one such equilibrium. For the weighted presentation, a complete characterization of the networks with this property is given. The necessary and sufficient condition is that the network has at most three routes that do traverse any edge in opposite directions, or it consists of several such networks connected in series. The corresponding problem for player-specific costs remains open. *Keywords:* Congestion games, network topology, existence of equilibrium.

1 Introduction

An exact potential for a noncooperative game is a function P on strategy profiles that exactly reflects the players' incentives to change their strategies. Whenever a single player moves to a different strategy, his gain or loss is equal to the corresponding change in P . In a game with a finite number of players and strategies, the existence of an exact potential implies that any improvement path, or chain of beneficial moves, must be finite: at some point, a (pure-strategy Nash) equilibrium is reached. Monderer and Shapley [32] showed that a finite game admits an exact potential if and only if it can be presented as a congestion game [37]. In this presentation, the players share a finite set E of resources, but may differ in which resources they are allowed to use. Specifically, each strategy of each player corresponds to a particular nonempty subset of E . The player's payoff from using the strategy is equal to the negative of the total cost of using the corresponding resources. The cost of each resource depends only on its identity and the number of users. It does not necessarily increase with congestion, and it may be negative (and equal to the negative of the gain from using the resource).

Restricting or expanding the meaning of 'congestion game' has a similar effect on the class of presentable finite games. A particularly natural restriction is increasing or (at least) nondecreasing cost functions: congestion never makes users better off. Two other possible restrictions are: singleton congestion games, where each strategy includes a single resource, and network congestion games, where resources are represented by edges in an undirected¹ graph and strategies correspond to routes, which are paths in the graph that connect the

¹ Directionality is viewed here as a limitation on the allowed usage of the network, rather than as an intrinsic property. See [27, 28].

player's origin and destination vertices. (The former restriction is a special case of the latter. It corresponds to a parallel network, which is one with only two vertices.) Examples of extensions are congestion games with player-specific costs (or payoffs [25]), in which players are differently affected by congestion, and weighted congestion games, in which their contributions to it (the players' "congestion impacts") differ. Different subsets of these alterations correspond to games with qualitatively different properties. A singleton congestion game with player-specific costs does not necessarily have a (pure-strategy) equilibrium if congestion makes players better, rather than worse off [21, 26]. By contrast, in the diametrically opposite case of nondecreasing cost function (i.e., a crowding game), at least one equilibrium always exists, although infinite improvement paths are possible [25, 26]. The existence of equilibrium in a singleton congestion game with nondecreasing cost functions is guaranteed also if the players differ in their weights rather than cost functions, but not if they differ in both respects [25]. This is not the case for network congestion games (on non-parallel networks), which may have no equilibrium even if the players differ *only* in their weights [14, 23] or only in their cost functions [20] and have identical allowable routes.

Libman and Orda [23] raised as an interesting subject for further research the problem of identifying all non-parallel networks for which the existence of equilibrium is guaranteed in all corresponding weighted congestion games. They added that series-parallel networks, which are built from single edges using only the operations of connecting networks in series or in parallel, may be especially interesting. For network congestion games with player-specific costs, Konishi [20] noted the similarity between the topological conditions for the existence of equilibrium and those for the *uniqueness* of the equilibrium in similar *nonatomic* games with a continuum of non-identical players, each with an infinitesimal congestion impact. (The existence of equilibrium in such games is not an issue since it is guaranteed by weak assumptions on the cost functions [39].) Specifically, a parallel network is a sufficient condition in both cases. The problem that these authors point to is thus the identification of all networks with the *topological (equilibrium) existence property* for a particular variety of network congestion games, which means that every game of that kind on the network has at least one equilibrium.

The topological existence property is particularly interesting for varieties of network congestion games which are capable of representing *all* finite games. As pointed out by Monderer [31], this is the case for network congestion games with player-specific costs. It is also the case for weighted network congestion games that are expanded to allow cost functions without self effect. In such games, the cost of a resource for a player may be a function of the total weight of the *other* users (see Section 2.3). For both kinds of network congestion games, it suffices to consider two-terminal (or single-commodity) networks, which have a single pair of origin and destination vertices where all players' routes start and terminate, respectively. Since these two kinds of network congestion games can be used to present any finite game, they cannot possibly have any special properties. Their significance lies in the information the presentation provides about the presented game. In particular, an equilibrium exists in every finite game that can be presented as a network congestion game on a (two-terminal) network with the topological existence property. This paper presents the solution to the problem of identifying all two-terminal networks with the topological

existence property for weighted network congestion games (expanded as indicated above). It also summarizes the known facts about the corresponding problem for network congestion games with player-specific costs and some additional related models.

1.1 Other properties of games related to the network topology

The topological equilibrium existence problem is substantially different from that of identifying classes of *cost functions* for which an equilibrium is guaranteed to exist. An example of such a class is linear (or affine) functions, that is, cost functions of the form $ax + b$, with $a, b \geq 0$. An equilibrium exists in every weighted network congestion game with linear cost functions, regardless of the network topology [14]. Linearity of the cost functions moreover implies that the game has a weighted potential (which changes *proportionally* to each player's gain or loss whenever only that player changes his strategy [32]), and this is the case also if the constant terms (the b 's) are player-specific (as well as resource-specific) [24]. Similarly, for network congestion games with player-specific costs but identical weights, a sufficient condition for the existence of equilibrium is that the players' (possibly, nonlinear) cost functions are identical up to additive (player- as well as resource-specific) constants [12, 21]. Indeed, such identity implies that the game has an exact potential. An immediate corollary of the last result is that every *singleton* congestion game with player-specific cost functions that are identical up to *multiplicative* constants (for example, linear cost functions without constant terms [16, 17]) has an equilibrium, and moreover has an ordinal potential (which changes in the same *direction* as the cost for of any single player who unilaterally changes his strategy [32]). Identity up to constants of the cost functions does not guarantee existence of equilibrium in singleton congestion games where the players differ also in their weights [30].

A similar distinction between the influences of the network topology and of the functional form of the cost functions also applies to the questions of the price of anarchy, the Pareto efficiency of the equilibria and the uniqueness of the equilibrium costs. Some of the known results concerning these properties of the equilibria are described below. In this account, 'network congestion game' without qualifiers refers to the lowest common denominator of identical players (that is, identical cost functions, weights and allowable routes, and in particular origin and destination vertices) and cost functions that are positive and increasing.

The (pure) price of anarchy [22, 35] in a game refers to some measure of social cost such as the maximum over all used strategies or the total (i.e., aggregate) cost. It is defined as the ratio between the social cost at the worst (pure-strategy) equilibrium in the game and the cost at the social optimum. The price of anarchy for a class of games is defined as the supremum over all games in the class. For any two-terminal network, the price of anarchy with respect to the maximum cost for network congestion games with (identical players and) linear cost functions on the network does not exceed $5/2$ [5]. For some networks, the price is lower. In particular [9], it is equal to 1 if and only if the network is an extension-parallel one [18], or a network with linearly independent routes [28], meaning that each (undirected) route has an edge that is not in any other route. The price of anarchy with respect to the total cost also depends on the network topology. It is $4/3$ for the class of network congestion games with linear cost functions on networks with linearly independent routes, but higher even for series-parallel networks [13]. Interestingly, the network topology

becomes essentially irrelevant for the price of anarchy with respect to the total cost (but not with respect to the maximum cost [9]) when the players' allowable routes may differ. For network congestion games with linear cost functions and player-specific allowable routes, the price is $5/2$ both for general networks and in the special case of parallel networks [3, 4, 5]. If the players may differ also in their weights, the price of anarchy with respect to the total cost rises to $(3 + \sqrt{5})/2 (\approx 2.618)$, but this again applies to both general and parallel networks [3, 4]. The irrelevance of the network topology also extends to nonatomic network congestion games, where, with linear cost functions and player-specific allowable routes, the price of anarchy with respect to the total cost is $4/3$ [38]. This maximum is already achieved in a game on a parallel network with only two edges. Whether or not the network topology is relevant for the intermediate model of splittable flow, in which the number of players is finite but they can split their flow arbitrarily among multiple routes, seems to be unknown [6].

An extension-parallel network is a necessary and sufficient condition also for weak Pareto efficiency of all equilibria in all (finite) network congestion games on the network, meaning that it is never possible to alter the players' equilibrium route choices in a way that benefits them all [18].² This result extends to games in which players may differ in their allowable routes, but have identical origin and destination vertices. It does not extend to games with player-specific costs, where inefficient equilibria may exist even with two-edge parallel networks. For nonatomic network congestion games, the network topology is relevant to efficiency both with identical and with player-specific costs. In both cases, a necessary and sufficient condition for weak Pareto efficiency of all equilibria in all such games on a two-terminal network is that the routes in the network are linearly independent, which is essentially the same condition as in the finite, identical-costs case (except that linear independence refers to undirected routes) [28]. In a sense, inefficient equilibria only occur in three particular two-terminal networks, which are the minimal such networks without the property of linearly independent routes. Linear independence of the routes is also a necessary and sufficient condition for the non-occurrence of Braess's paradox in all nonatomic network congestion games with player-specific costs on the network. For identical-costs games, this topological condition is replaced by the weaker condition of a series-parallel network [28].

The problem of the topological uniqueness of the equilibrium costs is relevant only for nonatomic network congestion games with player-specific costs. (In a game with identical cost functions, the equilibrium costs are always unique, and with a finite number of players, it is virtually impossible to guarantee uniqueness.) The class of two-terminal networks with guaranteed uniqueness of the players' equilibrium costs is characterized by five simple kinds of networks called the nearly parallel networks [27]. The complementary class of all two-terminal networks for which multiple equilibrium costs are possible consists of all the networks in which one of four particular "forbidden" networks is embedded. Similar results hold for network congestion games with finitely many players in which flow is splittable [36].

² An extension-parallel network moreover guarantees that all equilibria are strong [18]. This is 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.

Two additional issues related to the equilibrium existence problem in (finite) network congestion games are the efficient computation of equilibrium and the convergence to equilibrium of certain simple algorithms in which players sequentially choose best (or better) response strategies (see, e.g., [10, 11, 13, 19, 23, 25]). An example of such an algorithm is greedy best response: the players enter the game one after the other, and each new entrant chooses a best response to the strategies of the present players. This algorithm always reaches an equilibrium in an (unweighted) network congestion game on a two-terminal series-parallel network, but may fail for networks that are not series-parallel even though an equilibrium always exists [15]. For some additional results concerning complexity and convergence, see Section 4.

2 Preliminaries

2.1 Game theory

A *finite* noncooperative game Γ has a finite number n of players whose strategy sets are finite. A strategy profile $s = (s^1, s^2, \dots, s^n)$ in Γ , which specifies a strategy s^i to each player i , is a pure-strategy Nash equilibrium, or simply *equilibrium*, if none of the players can increase his payoff by unilaterally switching to another strategy.

Two games Γ and Γ' with identical sets of players are *isomorphic* [32] if for each player i there is a one-to-one correspondence between i 's strategy sets in Γ and Γ' , such that each strategy profile s in Γ yields the same payoffs to the players as the corresponding strategy profile s' in Γ' . Essentially, isomorphic games are just alternative presentations of a single game.

Two games Γ and Γ' with identical sets of players and respective strategy sets are *similar* if for each player the difference between the payoffs in Γ and Γ' can be expressed as a function of the other players' strategies. Equivalently, the gain or loss for a player from unilaterally switching from one strategy to another is always the same in both games. Similarity implies, in particular, that the two games are *best-response equivalent* [32, 33], 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. It follows that similar games have identical sets of equilibria.

A game Γ is an *exact potential game* [32] if it is similar to some game Γ' in which all players have the same payoff function. The players' common payoff function in Γ' 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* [32], 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*: 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*. This property differs from the finite improvement property in only requiring finiteness of best-(reply) improvement paths (where each new strategy is a best response for the moving player, who could not gain more by choosing some other strategy instead) 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 in which each player has to choose one of his strategies in each of the m games and the payoff is the sum of the resulting m payoffs [34]. Thus, the m games are played simultaneously but independently. Clearly, 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, u and v , 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 . 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 precedes e in the path to the vertex that follows it.

A *two-terminal network*, or simply *network*, G is an undirected multigraph together with a distinguished ordered pair of (distinct) *terminal* vertices, the *origin* o and the *destination* d , such that each vertex and each edge belongs to at least one path in which the first vertex is o and the last vertex is d . Any path with these first and last vertices will be called a *route* in G . A route can itself be viewed as a network. Specifically, 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.

The sub-network relation is a special case of the following one. 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 following operations any number of times in any order (see Figure 1):

- (a) The *subdivision of an edge*: its replacement by two edges with a single common end vertex.
- (b) The *addition of a new edge* joining two existing vertices.

³ This notion of embedding, which was introduced in [27], is more inclusive than that used in [28]. The latter only allows one kind of terminal division (see below), namely terminal extension. The difference between the two notions of embedding is roughly similar to that between a minor of a graph and a topological minor (see [7]).

- (c) The *subdivision of a terminal vertex*: addition of a new edge e that joins o or d with a new vertex v , followed by the replacement of the former by the latter as the end vertex in two or more edges originally incident with the terminal vertex.

Two networks are *homeomorphic* if they can be obtained from the same network by successive subdivision of edges. This relation represents a high degree of similarity between the networks: each can be obtained from the other by the insertion and removal of non-terminal vertices of degree two (which are incident with only two edges). Two networks will be identified if (they are *isomorphic* in the sense that) there is a one-to-one correspondence between the two sets of vertices, and another such correspondence between the sets of edges, such that (i) the incidence relation is preserved and (ii) the origin and destination in one network are paired with the origin and destination, respectively, in the other network.

Two networks G and G' may be connected *in parallel* if they have the same origin and the same destination but no other common vertices or edges, and *in series* if they have only one common vertex which is the destination in G and the origin in G' . In both cases, the set of vertices and the set of edges in the resulting network are the unions of the corresponding sets in G and G' , and the origin and destination are those in G and G' , respectively (as well as in G' and G , respectively, in the case of connection in parallel). The connection of an arbitrary number of networks in parallel or in series is defined recursively. Each of the connected networks is embedded in the wide sense in the network resulting from their connection.

A *parallel network* is a network that only has one edge or is made of several single-edge networks connected in parallel. A network is *nearly parallel* [27] if (i) it has only one route or (ii) it is made of two single-route networks connected in parallel to which any number of edges with identical end vertices were added, and each edge subdivided any number of times. Thus, depending on whether at most one or more than one edge was added, a nearly parallel network is homeomorphic to one of those in Figure 3 or Figure 5, respectively.




Figure 1. Embedding. The left 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.




Figure 2. Two-player weighted network congestion games (top) and their strategic (or normal) form (bottom). Dotted, dashed and solid edges are allowable to player 1, player 2 and both players, respectively. The allowable directions are indicated where needed. The players' weights are $w^1 = 2$ and $w^2 = 1$. All relevant costs other than those shown are zero.

2.3 Network congestion games

A *weighted network congestion game* on a (two-terminal⁴) network G is a finite, n -player game that is defined as follows. First, allowable direction and (possibly, empty) set of users are specified for each edge in G , such that (i) each edge is traversed in the allowed direction by at least one route and (ii) each player i has at least one *allowable route*, that is, a route r in G that includes only edges that i is allowed to use and traverses them in the allowed direction. The *strategy set* of each player i is the collection \mathcal{R}_i of his allowable routes. Second, a *weight* $w^i > 0$ is specified for each player i , which represents the player's congestion impact and is also assumed to be (weakly) connected with the cardinality of his strategy set: For all i and j with $w^i > w^j$, $|\mathcal{R}_i| \leq |\mathcal{R}_j|$.⁵ The total weight of the players whose chosen route includes an edge e , which is denoted by f_e , is the *flow* (or load) on e . The *cost* of using e for each player i , which may be positive or negative, is affected by the flow. The effect may take several forms, as detailed below. For each player i , the cost of an (allowable) route r is the sum of the costs of its edges. The player's *payoff* is the negative of the cost of his chosen route.

In this paper, the cost of an edge e for a player i may or may not involve *self effect*.⁶ That is, it may be a function of the total weight of all the users of e , which is the flow f_e , or only of

⁴ The assumption of a single origin–destination pair may be viewed as a normalization. Any 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 terminal vertex is incident with a single allowable edge for each player, which joins it with the player's corresponding terminal vertex in the original game.

⁵ The cardinality assumption is used only in the proof of Proposition 2. Whether or not it can be dispensed with I do not know. Doing so might strengthen some of the results presented below but weaken a little bit some of the others.

⁶ The special, and more familiar, case of mandatory self effect is considered in Section 4.3.

the total weight of the *other* users, $f_e - w^i$. Specifically, each edge e is associated with a pair of nondecreasing *cost functions* $c_e: (0, \infty) \rightarrow (-\infty, \infty)$ and $d_e: [0, \infty) \rightarrow (-\infty, \infty)$, such that its cost for each player i is given by

$$c_e(f_e) + d_e(f_e - w^i).$$

Lack of self effect is inconsequential in the special case of an *unweighted* network congestion game, where all weights are 1. In such a game, $d_e = 0$ can be assumed without loss of generality. However, if players do differ in their weights, then a non-zero d_e may mean that the cost of e is not the same for all players. An (*unweighted*) *network congestion game with player-specific costs* is a variant of the above model which extends this possibility by allowing the cost functions of different players to take arbitrarily different functional forms, but on the other hand, assumes that all weights are 1. In such a game, the cost of an edge e for a player i is $c_e^i(f_e)$, where $c_e^i: (0, \infty) \rightarrow (-\infty, \infty)$ is the corresponding nondecreasing cost function and (the flow) f_e is the total number of players using e .

2.4 Presentation theorem

The definitions of network congestion games involve rather specific structures. However, it turns out that the games themselves have no special properties. In fact, as the following theorem shows, *every* finite game can be presented as a weighted network congestion game *and* as a network congestion game with player-specific costs. Thus, the presentation only has to involve players that differ in their weights or players that differ in their cost functions. The existence of a presentation of the latter kind for every finite game was first pointed out by Monderer [31].

Theorem 1. Every finite game Γ is isomorphic both to a weighted network congestion game Γ' and to a network congestion game with player-specific costs Γ'' .

Proof. Suppose that the number n of players in Γ and the cardinality m of the largest strategy set are both at least two (otherwise the assertion is trivial), and that, for $1 \leq i < j \leq n$, player i 's number of strategies $m(i)$ is not greater than that of j (otherwise take 'player 1', 'player 2', ... below to mean the player with the smallest number of strategies, the second-smallest number, and so on). Index the strategies of each player i from 1 to $m(i)$. The indexing identifies each strategy profile with an element of M^n , where $M = \{1, 2, \dots, m\}$. Order all elements of M^n in the following way:

$$(1, 1, \dots, 1), (2, 2, \dots, 2), \dots, (m, m, \dots, m), \dots, (1, 2, \dots, 2), (2, 3, \dots, 3), \dots, (m, 1, \dots, 1), \quad (1)$$

where the order of the $m^n - 2m$ elements represented by in the middle ellipsis mark is arbitrary. With each element $s = (s^1, s^2, \dots, s^n)$ of M^n (which may or may not represent an actual strategy profile – the latter holds if $s^i > m(i)$ for some player i) associate two vertices u_s and v_s and an edge e_s joining them. The edge will be directed from u_s to v_s and be allowable to all players. Next, for each player i and integer $1 \leq k \leq m(i)$, consider all $s \in M^n$ with $s^i = k$ and list them according to their order in (1). For each pair s and t of successive entries in this list, add an edge that joins v_s and u_t , is directed from v_s to u_t and is allowable to player i only. Finally, identify all vertices of the form u_s , where s is one of the first m elements in (1), and denote this single vertex by o . Do the same for all vertices of the

form v_s , where s is one of the last m elements in (1), and denote the result by d . These terminal vertices, together with the other vertices and edges specified above, constitute a network G (see the example in Figure 2(b)), in which each allowable route r for each player i corresponds to some strategy s^i of i . Specifically, r includes all m^{n-1} edges e_t with $t^i = s^i$, alternating with $m^{n-1} - 1$ edges that are allowable to player i only. Different allowable routes to a player have no shared edges, and their only shared vertices are the terminal ones.

To complete the definitions of the weighted network congestion game Γ' and the network congestion game with player-specific costs Γ'' , which are both defined on G , it remains to specify the weights in the former and the cost functions in both games. The weight of player i ($= 1, 2, \dots, n$) is $w^i = 2n - i - 1$. These weights guarantee that every set of $n - 1$ players has a greater total weight than every set of $n - 2$ or fewer players. The former is equal to $(3n - 1)(0.5n - 1) + i$, where i is the unique player not in the set, and the latter is at most $(3n - 1)(0.5n - 1)$. The cost functions in Γ' will be defined as follows. For each edge in G of the form e_s , with $s \in M^n$ that is an actual strategy profile in Γ , c_{e_s} is identically zero and d_{e_s} is any nondecreasing function with

$$\begin{aligned} d_{e_s}(x) &= 0, & x &\leq (3n - 1)(0.5n - 1) \\ d_{e_s}((3n - 1)(0.5n - 1) + i) &= iK - h^i(s), & i &= 1, 2, \dots, n, \end{aligned} \quad (2)$$

where K is some number that is large enough to make d_{e_s} monotonic and h^i is player i 's payoff function in Γ . For $s \in M^n$ that is not an actual strategy profile, both c_{e_s} and d_{e_s} are identically zero. To offset the term iK in (2), all edges e that are allowable only to player i will have the same, constant cost functions $c_e = 0$ and $d_e = -iK/(m^{n-1} - 1)$. As explained above, strategy profiles in Γ are in a one-to-one correspondence with allowable route choices in G . The routes that correspond to a strategy profile s are such that exactly one edge, namely e_s , is used by all n players. Therefore, for each player i , only e_s and edges that are allowable only to i make a nonzero contribution to the player's cost, which by (2) is equal to $-h^i(s)$. The player's payoff, which is $h^i(s)$, is therefore the same as that in Γ .

The definition of Γ'' is similar, but simpler. For each strategy profile s ,

$$\begin{aligned} c_{e_s}^i(x) &= 0, & x &\leq n - 1 \\ c_{e_s}^i(n) &= K - h^i(s), & i &= 1, 2, \dots, n, \end{aligned}$$

where K is some sufficiently large number and h^i is player i 's payoff function in Γ . For each edge e allowable only to one player i , $c_e^i = -K/(m^{n-1} - 1)$. ■

A finite game Γ obviously has as more than a single pair of presentations as in Theorem 1. The "canonical" games Γ' and Γ'' constructed in the proof, which share the same network G , are just one such pair. Other presentations may be preferable in that certain properties of Γ are more easily inferable from them. An example of such an alternative presentation is shown in Figure 2(a). The 2×2 game in that example is presented as a weighted network congestion game on a particular five-edge network, the *Wheatstone network*. This is not the network G constructed in the proof of Theorem 1, which, for all 2×2 games, is the network in Figure 2(b). It can be immediately seen that the 2×2 game in (a) has two equilibria.

However, the existence of equilibrium in that game could be inferred without knowing the payoff matrices, the fact that the game is (essentially) symmetric, or even the number of players. As the next section shows, an equilibrium exists in *any* finite game that can be presented as a weighted network congestion game on the Wheatstone network. This is not the case for the network in Figure 2(b) (on which all 2×2 games are representable).

3 The Topological Existence Property

A network G has the *topological (equilibrium) existence property* for weighted network congestion games or for network congestion games with player-specific costs if every game of the specified kind on G has at least one (pure-strategy Nash) equilibrium. In view of Theorem 1, this means that every finite game that can be presented as such a network congestion game is guaranteed to have an equilibrium.

A sufficient condition for a network G to have the topological existence property for a particular kind of network congestion games is that G is embedded in the wide sense in a network that has that property. This is because any game on G that does *not* have an equilibrium can be “extended” to a game without an equilibrium on any network that is obtained from G by applying any of the three operations that define embedding in the wide sense (Figure 1). For example, the operation of adding a new edge can be made inconsequential by not allowing any player to use the edge, and the edge e that is created by terminal subdivision should be allowable with zero cost to all players.

Another sufficient condition for the topological existence property is that the network G is made of several networks with that property that are connected in series. The reason these networks bestow the topological existence property on G is that, as the proof of the following proposition shows, any network congestion game on G is the superposition (see Section 2.1) of such games on them.

Proposition 1. A two-terminal network made of two or more networks connected in series has the topological 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 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, every weighted network congestion game Γ on G can be presented as the superposition of m such games – one on each constituent network – and the same is true for a network congestion game with player-specific costs. In each of the m games, the players and their weights, as well as the cost functions and the allowable direction and players 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 an equilibrium, or this is so for every network congestion game with player-specific costs, then G also has the same property.

Conversely, if there is a weighted network congestion game without an equilibrium on G_k , for some $1 \leq k \leq m$, or there is some such network congestion game with player-specific costs, then a game with similar properties exists on G . For example, the superposition of the

game on G_k and games with zero costs on the other $m - 1$ networks is (isomorphic to) a game on G that does not have an equilibrium. ■

The rest of this section is concerned with weighted network congestion games, for which a complete characterization of the networks with the topological existence property is given.

3.1 Networks with the topological existence property

The simplest kind of network with the topological existence property for weighted network congestion games is a parallel network with no more than three edges (Figure 3(a)).

Proposition 2. Every weighted network congestion game Γ on a parallel network G with three or fewer edges has an 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 (precedes) edge e' if the latter is the first edge encountered with when moving along the cycle from e in the clockwise (respectively, counterclockwise) direction. The first part of the proof establishes the existence of an equilibrium under the additional assumption that no player has more than two allowable edges. The second part covers the general case. Both parts use the following simple result.

Claim 1. Let e and e' be two edges in G that are allowable to two players i and j . If both players use e , but only j would benefit from unilaterally moving to e' , then $w^j < w^i$.

The premise in Claim 1 means that the flows on e and e' are such that

$$\begin{aligned} (c_{e'}(f_{e'} + w^j) + d_{e'}(f_{e'})) - (c_e(f_e) + d_e(f_e - w^j)) &< 0 \\ &\leq (c_{e'}(f_{e'} + w^i) + d_{e'}(f_{e'})) - (c_e(f_e) + d_e(f_e - w^i)). \end{aligned}$$

The conclusion follows from the monotonicity of the cost functions $c_{e'}$ and d_e .

First part of the proof. Suppose that 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 \hat{w} of the players whose edge follows another allowable edge. There is obviously a unique strategy with $\hat{w} = 0$, which trivially satisfies the following:

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

Since \hat{w} cannot be greater than the total weight of the players, to prove that an equilibrium exists it suffices to establish the following.

Claim 2. For every strategy profile satisfying Q that is not an equilibrium, there is another strategy profile satisfying Q with a higher \hat{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 i 's edge e' . Such a move creates a strategy profile with a lower flow on e' and a

higher flow on e . That strategy profile may or may not have property Q . However, due to the monotonicity of the costs, 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. In that case, move the highest-weight such player from e to e' , and repeat doing that until no more players can benefit from that move. Necessarily, player i is not one of the movers. Indeed, i 's incentive to return to e' can only get lower with each move, and therefore Claim 1 implies that $w^j < w^i$ 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 some (possibly, empty) set of other players. The total weight w' of the players in that set must satisfy $w' < w^i$. Otherwise, for each of these players j , the monotonicity of the cost functions and the fact that $w^j < w^i$ would imply the following inequality:

$$\begin{aligned} & (c_{e'}(f_{e'}) + d_{e'}(f_{e'} - w^j)) - (c_e(f_e + w^j) + d_e(f_e)) \\ & \geq (c_{e'}(f_{e'} - w' + w^i) + d_{e'}(f_{e'} - w')) \\ & \quad - (c_e(f_e + w') + d_e(f_e + w' - w^i)). \end{aligned}$$

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 (strictly) positive since it gives the reduction in the cost to i when he moved from e' to e . This shows that the above inequality, and hence also $w' \geq w^i$, actually cannot hold.

The result that $w^i - w'$ is positive means that f_e is higher, and $f_{e'}$ is lower, than the respective flows in the original strategy profile. Therefore, there are still no players who would gain from moving to e from the third edge in G or from moving to that edge from e' . Hence, Q holds for both the original and the new strategy profiles. In the latter, the total weight \widehat{w} of the players whose edge follows another allowable edge is higher by $w^i - w'$ than in the former. This completes the proof of Claim 2.

Second part of the proof. Suppose that Γ has some players i with three allowable edges, possibly in addition to players j with only one or two. Re-index the players in the game in such a way that, for some $k \geq 1$, the inequalities $j < k \leq i$ hold for all players i and j as above, who differ in their number of strategies, and $w^j \geq w^i$ holds for *all* i and j with $j < i$. (The cardinality assumption in the definition of network congestion game implies that such re-indexing is possible.) For each player i , define Γ_i as the game obtained from Γ by ‘‘taking out’’ i and all the higher-index players, so that these players do not choose routes and do not contribute to the flows. For example, Γ_k is the game in which only the players in Γ with one or two allowable edges participate. This game may actually have no players.

It follows from the first part of the proof that Γ_k has an equilibrium. To prove that an equilibrium exists also in every Γ_i with $i > k$ (and hence in Γ), it suffices to show that, for every such i , the existence of an equilibrium in Γ_{i-1} implies the same for Γ_i . In fact, for any equilibrium in Γ_{i-1} , simply choosing a best response strategy for player i gives an equilibrium in Γ_i . 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 $w^j \geq w^i$, and since moving from e to another edge e' is not beneficial to i , it follows from Claim 1 that the same applies to j . ■




Figure 3. Two-terminal networks with the topological existence property. Every weighted network congestion game on any of these networks has a (pure-strategy Nash) equilibrium. The two gray curves in the parallel network (a) are optional edges.

By Propositions 1 and 2, any network that can be constructed by connecting in series two or more parallel networks as in Figure 3(a) (for example, the figure-eight network) has the topological existence property. The next proposition shows that the property also holds for the networks in Figure 3(b)–(e), which cannot be constructed in this way. Indeed, the Wheatstone network (Figure 3(e)) is not even series-parallel, meaning that it cannot be constructed from networks with single edges by *any* sequence of operations of connecting networks in series or in parallel.

Proposition 3. Every weighted network congestion game Γ on any of the networks in Figure 3 has an equilibrium.

This result is an immediate corollary of Proposition 2 and the following lemma. The lemma shows that every network congestion game Γ as in Proposition 3 is similar (see Section 2.1) to a game on a particular parallel network. That game is obtained from Γ by a procedure (“parallelization”) that involves transformation of 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 forms are intimately connected.

Lemma 1. Every weighted network congestion game Γ on any of the networks G in Figure 3 is similar to such a game $\tilde{\Gamma}$ on a parallel network with three edges.

Proof. Suppose that G is one of the networks in Figure 3(b)–(e). (For (a), the assertion is trivial.) Let \tilde{G} be the parallel network obtained from G by *contracting* [7] edges e' and e'' (or only the former, if G has only four edges), that is, performing the (one-sided) inverse of terminal subdivision (Figure 1(c)), which eliminates the edge and its non-terminal vertex. Each of the three routes in \tilde{G} corresponds to a route in G , in that the former’s single edge is the unique edge in the latter that did not undergo contraction. This correspondence between routes is one-to-one and onto, with one exception. The single exception is route $e_4 e_1 e_3$ in the Wheatstone network (e), which does not have a corresponding route in the parallel network. The omission of that route is inconsequential since, by symmetry, it suffices to consider network congestion games on the Wheatstone network in which the allowable direction of e_1 is from u to v . Thus, it suffices to consider weighted network

congestion games Γ on G in which every route r that is allowable for some player includes a unique edge that also belongs to the corresponding parallel network \tilde{G} . The next step is to describe the corresponding weighted network congestion game $\tilde{\Gamma}$ on \tilde{G} .

The following description concerns the case in which G is the Wheatstone network (e), so that \tilde{G} is the parallel network with edges e_1, e_3 and e_4 . The other three cases ((b)–(d)) are similar (actually, simpler). The game $\tilde{\Gamma}$ on \tilde{G} inherits its set of players, their weights and the strategy sets from the game Γ on G (with the identification of routes in G and \tilde{G} described above). The cost functions in $\tilde{\Gamma}$ (which are marked by a tilde) are derived from those in Γ (without a tilde) as follows. For $0 \leq y < x \leq w$, where $w = \sum_i w^i$ is the players' total weight,

$$\begin{aligned} \tilde{c}_{e_1}(x) &= c_{e_1}(x), & \tilde{d}_{e_1}(y) &= d_{e_1}(y), \\ \tilde{c}_{e_3}(x) &= c_{e_3}(x) - d_{e_3}''(w - x), & \tilde{d}_{e_3}(y) &= d_{e_3}(y) - c_{e_3}''(w - y), \\ \tilde{c}_{e_4}(x) &= c_{e_4}(x) - d_{e_4}'(w - x), & \tilde{d}_{e_4}(y) &= d_{e_4}(y) - c_{e_4}'(w - y). \end{aligned}$$

It remains to show that the games Γ and $\tilde{\Gamma}$ are similar. That is, for every player i , the difference between the costs to i in Γ and $\tilde{\Gamma}$ can be expressed as a function of the route choices of the other players. If i 's route includes e_1 (and, thus, does not include e_3 or e_4), the difference can be written as

$$c_{e_1}'(w_4^{-i} + w^i) + d_{e_1}'(w_4^{-i}) + c_{e_3}''(w_3^{-i} + w^i) + d_{e_3}''(w_3^{-i}), \quad (3)$$

where w_j^{-i} is the total weight of the players other than i whose route does *not* include e_j . The same expression gives the difference between the costs in Γ and $\tilde{\Gamma}$ also if i 's route r does include either e_3 or e_4 . For example, if r includes (only) the former, the difference is

$$\begin{aligned} &(c_{e_1}'(w - f_{e_4}) + d_{e_1}'(w - f_{e_4} - w^i) + c_{e_3}(f_{e_3}) + d_{e_3}(f_{e_3} - w^i)) \\ &\quad - (\tilde{c}_{e_3}(f_{e_3}) + \tilde{d}_{e_3}(f_{e_3} - w^i)), \end{aligned}$$

which equals (3). Thus, the difference is independent of i 's route, as had to be shown. ■

Parenthetically, the assertion of Lemma 1 cannot be strengthened to isomorphism between Γ and $\tilde{\Gamma}$. In other words, the class of weighted network congestion games on the networks in Figure 3 and the subclass obtained by only considering the parallel networks shown in (a) are not equal. For example, it is not difficult to see that the 2×2 game in Figure 2(a) cannot be presented as a weighted network congestion game on any parallel network.

3.2 Networks without the topological existence property

A network without the topological existence property can be obtained from a network homeomorphic to one of those in Figure 3 (with that proviso that the two optional edges in (a) actually exist) simply by the addition of a single edge (with existing end vertices). *Any* such addition will have that effect. This is because the resulting network necessarily has one or more of those in Figure 4 embedded in it in the wide sense. As the proof of the following proposition shows, there are three-player games on the networks in Figure 4(b)–(e) and a four-player game on the four-edge parallel network (a) which do not have an equilibrium. It can be shown that these numbers of players are minimal for non-existence of equilibrium. In

particular, every weighted network congestion game with three or fewer players on any parallel network has an equilibrium.

Proposition 4. A weighted network congestion game without an equilibrium exists on each of the networks in Figure 4.

Proof. The proof comprises the following four examples.

Example 1. Four players, with weights $w_1 = 1$, $w_2 = 2$ and $w_3 = w_4 = 3$, choose routes in the network in Figure 4(a). 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. Edge e_1 has the constant cost $c_{e_1} = 16$. The other edges have variable costs, with $c_{e_2}(1) = 2, c_{e_2}(3) = 3, c_{e_2}(4) = 15, c_{e_2}(5) = 17$; $c_{e_3}(1) = 4, c_{e_3}(3) = 10, c_{e_3}(4) = 14$ and $c_{e_4}(2) = 2, c_{e_4}(3) = 11, c_{e_4}(5) = 12$. In addition, $d_e = 0$ for all edges e except e_4 , for which $d_{e_4}(0) = 0, d_{e_4}(2) = 1$ and $d_{e_4}(3) = 6$. 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 former case player 2 can decrease his cost from 3 to 2 by (unilaterally) changing his choice to “right”, and in the latter case, he can decrease it from 18 to 17 by changing to “left”. This proves that an equilibrium does not exist.

Example 2. Three players, with weights $w_1 = 3$ and $w_2 = w_3 = 4$, choose routes in the network in Figure 4(b) or in that in (c). 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_3 is only allowable to player 3, who is not allowed to use e_4 . Thus, there are two allowable routes for each player: “left”, which includes e_5 , and “right”, which does not. The two “private” edges have constant costs: $c_{e_2} = 7$ and $c_{e_3} = 13$. The other edges have variable costs: $c_{e_1}(x) = x$, $c_{e_4}(x) = 0.75(0.25x - 2)^9 + 15$ and $c_{e_5}(x) = x$. For all edges e , $d_e = 0$. 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 an equilibrium does not exist.




Figure 4. Networks without the topological existence property. On each of these networks there is a weighted network congestion game that does not have a (pure-strategy Nash) equilibrium.

Example 3. Three players, with weights $w_1 = 1$ and $w_2 = w_3 = 2$, 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_3 is only allowable to player 3, who is not allowed to use e_4 . Thus, there are two allowable routes to each player: “left”, which does not include e_5 , and “left”, which does. The two “private” edges have constant costs: $c_{e_2} = 3$ and $c_{e_3} = 9$. The other edges have variable costs, with $c_{e_1}(1) = 1, c_{e_1}(2) = 2, c_{e_1}(3) = 8; c_{e_4}(1) = 2, c_{e_4}(2) = 10, c_{e_4}(3) = 12$ and $c_{e_5}(x) = 4x$. For all edges $e, d_e = 0$. 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 an equilibrium does not exist.

Example 4. Three players, with weights $w_1 = 1, w_2 = 5$ and $w_3 = 10$, choose routes in the network in Figure 4(e). The only restrictions are that edge e_2 is only allowable to player 2, who is not allowed to use e_1 , and e_3 is only allowable to player 3, who is not allowed to use e_4 . Thus, there are two allowable routes to each player: “left”, which does not include e_5 , and “left”, which does. Three of the edges have constant costs, $c_{e_2} = 1.3, c_{e_3} = 6.25$ and $c_{e_5} = 40$, and three have variable costs, $c_{e_1}(x) = 2x, c_{e_4}(x) = 5x$ and $c_{e_6}(x) = 3.55\sqrt{x}$. For all edges $e, d_e = 0$. 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 an equilibrium does not exist.

Another example of a game without an equilibrium on the network in Figure 4(e) can be obtained from Example 3 by simply setting $c_{e_6} = 0$. ■

3.3 Characterization

The main result in this section is that, for weighted network congestion games, the networks in Figure 3 and Figure 4 are in a sense the only networks with and without the topological existence property, respectively. It is based on a graph theoretic result [27, Proposition 2.1], which identifies one or more of the networks in these figures in every two-terminal network.

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

- (i) G is homeomorphic to one of the networks in Figure 3, or to a network made of several such networks connected in series.
- (ii) None of the networks in Figure 4 is embedded in the wide sense in G .
- (iii) G has the topological existence property.

Proof. Networks homeomorphic to one of those presented in Figure 3 are a special case of nearly parallel networks (see Section 2.2). They differ from the other nearly parallel networks, which are homeomorphic to one of those shown in Figure 5, in that the four-edge parallel network (Figure 4(a)) is not embedded in them in the wide sense. The networks in Figure 4(b)–(e) are not nearly parallel. They are called the *forbidden networks* in [27], where it is proved that one, and only one, of the following two conditions holds for every two-terminal network G :

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

If a network G satisfies (ii') but not (i'), then it does not satisfy (i) or (ii). It hence follows from Proposition 4 that G also does not satisfy (iii). If G satisfies (i') but not (ii'), there are two cases to consider. If (i) does not hold, then (since (i') does hold) the network in Figure 4(a) is embedded in the wide sense in G . Hence, (ii) does not hold and, by Example 1, the same is true for (iii). If (i) does hold, then it follows from Propositions 1 and 3 that (iii) also holds, which by Proposition 4 implies the same for (ii). This completes the proof of the equivalence of conditions (i), (ii) and (iii): either all of them hold, or none of them holds. ■

An additional, strikingly simple characterization of networks with the topological existence property follows as an immediate corollary from the following observation. The four or five routes in each of the networks in Figure 4 have the property that no two routes pass through any edge in opposite directions. Each of the operations that define embedding in the wide sense can obviously only increase the number of routes with that property or leave it unchanged. By contrast, the maximum number of such routes in each of the networks in Figure 3 is three. In view of Lemma 2, this implies that, to tell whether a given network G has the topological existence property, it suffices to record the maximum number of routes as above in the networks G is made of, in the sense of connection of networks in series. This proves the following.

Theorem 2. For a two-terminal network G that is not made of two or more such networks connected in series, a weighted network congestion game without an equilibrium exists on G if and only if there are four routes in the network such that no two routes pass through any edge in opposite directions.

It follows from Theorem 2 that there is, for example, a weighted network congestion game on the (underlying undirected) network in Figure 2(b) that does not have an equilibrium. This is of course also an immediate corollary of the result that every 2×2 game can be presented as such a weighted network congestion game (see the proof of Theorem 1).

One may wonder whether the non-existence in of four routes as in Theorem 2 in a network actually guarantees more than just the existence of equilibrium, that is, whether there are any stronger properties that are common to all weighted network congestion games on such networks. One such property might be the existence of a (generalized ordinal) potential (see Section 2.1). However, as the following example shows, this property is in fact not guaranteed. Even in a three-player game on a three-edge parallel network, improvement (and even best-improvement) cycles may exist. Although such a game always has an equilibrium, a specific order of moves may be required to get there.

Example 5. Three players, with weights $w_1 = w_2 = 1$ and $w_3 = 2$, choose routes in the parallel network with edges e_1, e_2, e_3 . Each player i can use all edges except e_i . The cost functions are: $c_{e_1}(x) = 16.75 - 9/x$, $c_{e_2}(x) = 3x + 6$, $c_{e_3}(x) = 8x$, $d_{e_1}(x) = x^2$, $d_{e_2}(x) = d_{e_3}(x) = 0$. It can be verified that, starting with the strategy profile in which players 1 and 2 use e_3 and player 3 uses e_2 , the following is a (best-) improvement cycle: player 1 moves to e_2 , player 2 moves to e_1 , player 3 moves to e_1 , player 1 moves to e_3 , player 2 moves to e_3 , and player 3 moves to e_2 , thus completing the cycle. Note that an equilibrium would be (immediately) reached if player 2 (rather than 1) moved first (to e_1), and a different equilibrium would be reached if player 3 (rather than 2) moved second (to e_1).

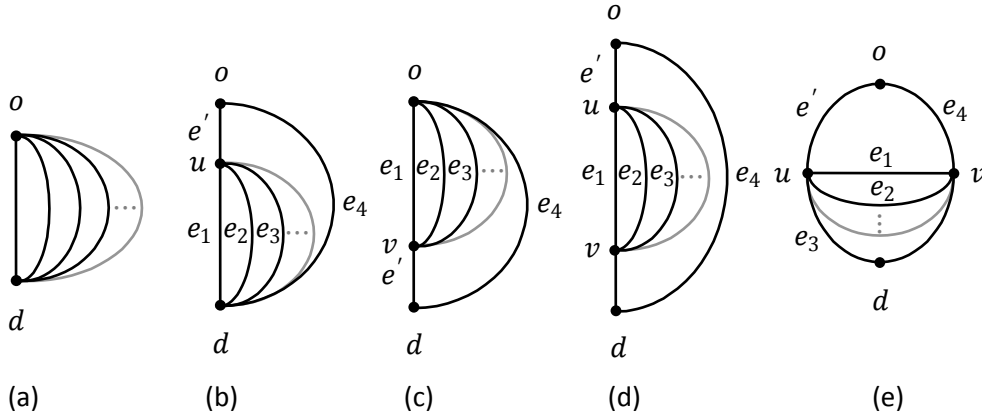


Figure 5. The parallel network (a) has the topological existence property for weighted network congestion games with mandatory self effect. The networks (b)–(e) do not have that property. Gray ellipsis mark and curve represent (any number of) optional edges.

4 Related Models and Open Problems

4.1 Player-specific costs

The topological existence property for network congestion games with player-specific costs is not equivalent to the corresponding property for weighted network congestion games (Section 3). Specifically, the former is less demanding: it holds not only for the networks G that satisfy condition (i) (or (ii)) in lemma 2 but also for certain other networks. In particular, an equilibrium exists in every network congestion games with player-specific costs on any parallel network, regardless of the number of edges [25]. By a parallelization argument similar to that in Lemma 1 it follows that the topological existence property actually holds for all nearly parallel networks that are homeomorphic to one of those in Figure 5 (a)–(d) [29].⁷

The main open problem regarding the topological existence property for network congestion games with player-specific costs is whether, and to what extent, the property holds for networks that are not nearly parallel or made of several nearly parallel networks connected in series. In particular, it is not known whether the forbidden networks (Figure 4(b)–(e)) have this property. A partial result is that each of the following single-edge additions gives a network without the topological existence property [29]: an edge with end vertices o and u in Figure 4(b) (equivalently, v and d in (c) or o and d in (d)), end vertices u and v in Figure 4(e), or end vertices o and d in the Wheatstone network (Figure 3(e)).

From a computational complexity point of view, finding an equilibrium in a network congestion game with player-specific costs on a parallel network (that is, a singleton congestion game), and hence also in such a game on each of the nearly parallel networks mentioned above, is not difficult. Starting with any strategy profile, there is a best-

⁷ The parallelization argument partially applies also to the remaining nearly parallel networks, which are those homeomorphic to a network as in Figure 5(e) [29]. The limitation is that the argument only applies to games in which all the allowable routes that include both u and v pass these vertices in the same order: either u first or v first.

improvement path that ends at an equilibrium, with a length that is polynomial in the number of players and strategies [25]. For a network congestion game with player-specific costs on a general network, it may be computationally difficult to determine whether an equilibrium exists. Ackermann and Skopalik [1] showed that this problem is in fact NP-complete even with only two players.

4.2 Resource-symmetric games

In a *resource-symmetric* (often referred to simply as “symmetric”) weighted network congestion game, all players share the same set of allowable edges, which without loss of generality may be assumed to include all edges. The topological existence property for such games is less demanding than in the general case considered in Section 3. In particular, it holds for all parallel networks. This can easily be proved constructively by employing the *greedy best response* algorithm [15] used in the second part of the proof of Proposition 2, that is, by letting the players enter the game one by one with heavier players entering first. As is the case for network congestion games with player-specific costs (Section 4.1), a parallelization argument extends this result to the nearly parallel networks in Figure 5(b)–(d) (and partially to (e); see footnote 7) [29]. However, whether every network that has the topological existence property for network congestion games with player-specific costs also has that property for resource-symmetric weighted network congestion games, or vice versa, is an open problem. For some of the examples (in Section 4.1) of networks without the first property, it is not known whether or not the second property holds, and for other networks, the reverse is true. An example of the latter is the network with linearly independent routes that is obtained from that in Figure 4(c) by subdividing e_1 and joining the resulting new vertex with o by a new edge. A resource-symmetric weighted network congestion games on this network that does not have an equilibrium is presented in [29].

4.3 Mandatory self effect

Another natural subclass of weighted network congestion games is obtained by mandating self effect: d_e must be identically zero for all edges e . Like resource symmetry, mandatory self effect adds to the set of networks with the topological existence property all parallel networks with more than three edges (Figure 5(a)). However, unlike resource symmetry, it adds essentially *only* these networks.

The reason why restriction to cost functions with self effect guarantees the existence of equilibrium in every weighted network congestion game on a parallel network is that it entails that the cost of an edge is the same for every player who is allowed to use it. This equality implies that the game has a (generalized ordinal) potential [10, 11]. The reason why that restriction does not extend the topological existence property to other nearly parallel networks with more than three routes is that games that satisfy the restriction but do not have an equilibrium exist on each of the networks in Figure 5(b)–(e). The game presented in the following example is very similar to that in Example 1. In fact, the latter can be obtained from the former by parallelization (see the proof of Lemma 1).

Example 6. 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 5(b)–(e). Each player has two allowable routes, each of which includes exactly one of the edges e_1 , e_2 , e_3 and e_4 . The “left” route for player 1, 2, 3 and 4

includes e_2 , e_2 , e_1 and e_3 , respectively, and the “right” route includes e_3 , e_4 , e_2 and e_4 , respectively. Edge e_1 has the constant cost $c_{e_1} = 16$. Edges e_2 , e_3 , e_4 and e' have variable costs, with $c_{e_2}(1) = 2, c_{e_2}(3) = 3, c_{e_2}(4) = 15, c_{e_2}(5) = 17; c_{e_3}(1) = 4, c_{e_3}(3) = 10, c_{e_3}(4) = 14; c_{e_4}(2) = 9, c_{e_4}(3) = 18, c_{e_4}(5) = 19$ and $c_{e'}(6) = 1, c_{e'}(7) = 6, c_{e'}(9) = 7$. For all other edges e (if the network has them), $c_e = 0$. In addition, $d_e = 0$ for all e . 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 former case player 2 can decrease his cost from 10 to 9 by (unilaterally) changing his choice to “right”, and in the latter case, he can decrease it from 19 to 18 by changing to “left”. This proves that an equilibrium does not exist.

Using arguments similar to those in the proof of Lemma 2, it is now not difficult to prove the following.

Theorem 3. For a two-terminal network G that is not made of two or more such networks connected in series, the following conditions are equivalent:

- (i) G is homeomorphic to a parallel network, or it has at most three routes that do not pass through any edge in opposite directions.
- (ii) An equilibrium exists in very weighted network congestion game on G in which all cost functions have self effect.

For networks that do not satisfy condition (i) in Theorem 3, it may be computationally difficult to decide whether an equilibrium exists in a given weighted network congestion game where all cost functions exhibit self effect. Dunkel and Schulz [8] showed, in fact, that without any assumptions on the network topology this decision problem is NP-complete even in the special cases of resource symmetry or only four players.

The equilibrium existence decision problem is NP-complete even in the case of parallel networks if the players also have different cost functions [8]. Thus, with player-specific weights *and* costs, the network topology is essentially irrelevant to the complexity of deciding whether an equilibrium exists. A game of this kind may have no equilibrium even on a three-edge parallel network with only three players [25].

4.4 Matroid congestion games

Every network topology entails a particular set of combinatorial restrictions on the players’ strategy sets in all corresponding network congestion games. For example, in every such game, different strategies are incomparable in that the set of edges in one strategy is not a subset of that in another. The restrictions take an extreme form in the case of parallel networks, which correspond to singleton congestion games. This observation leads to the question of whether the existence of equilibrium in this and similar classes of network congestion games can be linked directly to the combinatorial structure of a player’s strategy space, rather than to the network topology that gives rise to that structure. Specifically, Ackermann et al. [2] presented the following combinatorial version of the equilibrium existence problem: What is the most general combinatorial structure for which an

equilibrium is guaranteed to exist in every corresponding weighted congestion game with mandatory self effect, and what is that structure for player-specific costs?

As Ackermann et al. [2] showed, the most general games of both kinds are *matroid congestion games*, in which the strategy space of each player consists of the bases of a matroid on the set of resources. These games and singleton congestion games share the property (which reflects the corresponding property of bases of a matroid) that all strategies of a player include the same number of resources. However, they allow for much more varied and elaborate combinatorial structures, for example, strategy sets that consist of all *pairs* of resources. A noteworthy aspect of the results of Ackermann et al. is that they do not take into account how the strategy spaces of different players interweave. This means that the existence of equilibrium in weighted congestion games with mandatory self effect and in games with player-specific costs may be guaranteed even if some (or all) of the players have strategy spaces that do *not* consist of the bases of a matroid (for example, if some strategies of a player include fewer edges than others). The results only entails that, with such strategy sets, it is possible to systematically substitute a different edge for each allowable edge for each player, such that with the *modified* strategy sets the existence of equilibrium is not guaranteed.

The positive part of the solution to the combinatorial equilibrium existence problem does apply to network congestion games [2]. However, its usefulness for the graph-theoretic version studied in this paper is limited. This assertion is based on the following fact.

Proposition 5. In a network congestion game on a two-terminal network G , the strategy set of a player i consists of the bases of a matroid on the set of edges if and only if the sub-network of G that includes only i 's allowable edges is parallel or is made of several parallel networks connected in series.

Proof. It suffices to show that the first condition (the matroid property) is equivalent to the following graph theoretic one: all allowable routes for i have the exact same vertices and pass them in the same order. Since different routes have incomparable sets of edges, these sets of edges are the bases of a matroid if and only if they satisfy the bijective exchange axiom [40]: 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 gives a third (or the same) 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. ■

References

1. Ackermann, H. and Skopalik, A. (2007). On the complexity of pure Nash equilibria in player-specific network congestion games. Lecture Notes in Computer Science 4858, 419–430.
2. Ackermann, H., Röglin, H. and Vöcking, B. (2009). Pure Nash equilibria in player-specific and weighted congestion games. Theoretical Computer Science 410, 1552–1563.

3. Awerbuch, B., Azar, Y. and Epstein, A. (2005). The price of routing unsplittable flow. *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, 57–66.
4. Caragiannis, I., Flammini, M., Kaklamanis, C., Kanellopoulos, P. and Moscardelli, L. (2006). Tight bounds for selfish and greedy load balancing. *Lecture Notes in Computer Science* 4051, 311–322.
5. Christodoulou, G. and Koutsoupias, E. (2005). The price of anarchy of finite congestion games. *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, 67–73.
6. Cominetti, R., Correa, J. R. and Stier-Moses, N. E. (2009). The impact of oligopolistic competition in networks. *Operations Research*, in print.
7. Diestel, R. (2005). *Graph Theory*. Third edition. Graduate Texts in Mathematics, Vol. 173. Springer-Verlag, Berlin, Heidelberg, New York.
8. Dunkel, J. and Schulz, A. S. (2008). On the complexity of pure-strategy Nash equilibria in congestion and local-effect games. *Mathematics of Operations Research* 33, 851–868.
9. Epstein, A., Feldman, M. and Mansour, Y. (2009). Efficient graph topologies in network routing games. *Games and Economic Behavior* 66, 115–125.
10. Even-Dar, E., Kesselman, A. and Mansour, Y. (2003). Convergence time to Nash equilibria. *Lecture Notes in Computer Science* 2719, 502–513.
11. Fabrikant, A., Papadimitriou, C. and Talwar, K. (2004). The complexity of pure Nash equilibria. *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, 604–612.
12. Facchini, G., Van Meegen, F., Borm, P. and Tijs, S. (1997). Congestion models and weighted Bayesian potential games. *Theory and Decision* 42, 193–206.
13. Fotakis, D. (2009). Congestion games with linearly independent paths: Convergence time and price of anarchy. *Theory of Computing Systems*, in print.
14. Fotakis, D., Kontogiannis, S. and Spirakis, P. (2005). Selfish unsplittable flows. *Theoretical Computer Science* 348, 226–239.
15. Fotakis, D., Kontogiannis, S. and Spirakis, P. (2006). Symmetry in network congestion games: Pure equilibria and anarchy cost. *Lecture Notes in Computer Science* 3879, 161–175.
16. Gairing, M., Monien, B. and Tiemann, K. (2006). Routing (un-) splittable flow in games with player-specific linear latency functions. *Lecture Notes in Computer Science* 4051, 501–512.

17. Georgiou, C., Pavlides, T. and Philippou, A. (2006) Network uncertainty in selfish routing. Proceedings of the 20th IEEE International Parallel and Distributed Processing Symposium.
18. Holzman, R. and Law-yone (Lev-tov), N. (2003). Network structure and strong equilibrium in route selection games. *Mathematical Social Sciences* 46, 193–205.
19. Jeong, S., McGrew, R., Nudelman, E., Shoham, Y. and Sun, Q. (2005). Fast and compact: A simple class of congestion games. Proceedings of the 20th National Conference on Artificial Intelligence, 489–494.
20. Konishi, H. (2004). Uniqueness of user equilibrium in transportation networks with heterogeneous commuters. *Transportation Science* 38, 315–330.
21. Konishi, H., Le Breton, M. and Weber, S. (1997). Pure strategy Nash equilibrium in a group formation game with positive externalities. *Games and Economic Behavior* 21, 161–182.
22. Koutsoupias, E. and Papadimitriou, C. (2009). Worst-case equilibria. *Computer Science Review* 3, 65–69.
23. Libman, L. and Orda, A. (2001). Atomic resource sharing in noncooperative networks. *Telecommunication Systems* 17, 385–409.
24. Mavronicolas, M., Milchtaich, I., Monien, B. and Tiemann, K. (2007). Congestion games with player-specific constants. *Lecture Notes in Computer Science* 4708, 633–644.
25. Milchtaich, I. (1996). Congestion games with player-specific payoff functions. *Games and Economic Behavior* 13, 111–124.
26. Milchtaich, I. (1998). Crowding games are sequentially solvable. *International Journal of Game Theory* 27, 501–509.
27. Milchtaich, I. (2005). Topological conditions for uniqueness of equilibrium in networks. *Mathematics of Operations Research* 30, 225–244.
28. Milchtaich, I. (2006). Network topology and the efficiency of equilibrium. *Games and Economic Behavior* 57, 321–346.
29. Milchtaich, I. (2006). The equilibrium existence problem in finite network congestion games. *Lecture Notes in Computer Science* 4286, 87–98.
30. Milchtaich, I. (2009). Weighted congestion games with separable preferences. *Games and Economic Behavior* 67, 750–757.
31. Monderer, D. (2007). Multipotential games. Proceedings of the 20th International Joint Conference on Artificial intelligence, 1422–1427.

32. Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and Economic Behavior* 14, 124–143.
33. Morris, S. and Ui, T. (2004). Best response equivalence. *Games and Economic Behavior* 49, 260–287.
34. von Neumann, J. and Morgenstern, O. (1953). *Theory of Games and Economic Behavior*. Third edition. Princeton University Press, Princeton, NJ.
35. Papadimitriou, C. H. (2001). Algorithms, games, and the Internet. *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, 749–753.
36. Richman, O. and Shimkin, N. (2007). Topological uniqueness of the Nash equilibrium for atomic selfish routing. *Mathematics of Operations Research* 32, 215–232.
37. Rosenthal, R. W. (1973). A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 65–67.
38. Roughgarden, T. (2003). The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences* 67, 341–364.
39. Schmeidler, D. (1970). Equilibrium points of nonatomic games. *Journal of Statistical Physics* 7, 295–300.
40. White, N. (Ed.) (1986). *Theory of Matroids*. *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, UK.