

Network Characterizations for Excluding Braess's Paradox

Xujin Chen¹ · Zhuo Diao¹ · Xiaodong Hu¹

Published online: 14 October 2016
© Springer Science+Business Media New York 2016

Abstract Braess's paradox exposes a counterintuitive phenomenon that when travelers selfishly choose their routes in a network, removing links can improve the overall network performance. Under the model of nonatomic selfish routing, we characterize the topologies of k -commodity undirected and directed networks in which Braess's paradox never occurs. Our results strengthen Milchtaich's series-parallel characterization (Milchtaich, Games Econom. Behav. **57**(2), 321–346 (2006)) for the single-commodity undirected case.

Keywords Nonatomic selfish routing · Braess's paradox · Single-commodity network · Multcommodity network · Series-parallel graph

1 Introduction

A basic task of network management is routing traffic to achieve the best possible network performance, e.g., to minimize the maximum latency. However, it is usually difficult or even impossible to implement centralized optimal routing in large network systems where a number of network users (players) are playing *selfish routing* games [8, 16, 23]. In these games, players choose routes in the networks for traveling from their origins to their destinations. These route choices are selfish in the

Research supported in part by NNSF of China under Grant No. 11531014 and 11222109.

✉ Xujin Chen
xchen@amss.ac.cn

¹ Institute of Applied Mathematics, Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing 100190, China

sense that the players only aim to minimize their *own* latencies. It is well known that Nash equilibria of selfish routing games might not be socially optimal [17]. A dazzling example is Braess's paradox [4], which exposes the seemingly counterintuitive phenomenon that less route options for the players lead to shorter travel time at the equilibrium – subnetworks have better performance under the selfish behaviors. The natural question arises as to which topologies of networks are immune to the inefficiency due to the occurrence of Braess's paradox. The characterization of network topologies, which model fixed infrastructures, is independent of the changeable latency functions and traffic demands. Once such a paradox-free network is established, no matter how the latency functions and traffic demands change, the entire network remains a good venue for selfish players. The goal of this paper is to characterize paradox-free network topologies for nonatomic players each routing a negligible portion of the overall traffic, where the nonatomic routing can be viewed as a mathematical idealization of a very large population of individuals.

1.1 Noatomic Selfish Routing

We are concerned with both undirected and directed networks, and model them by multigraph or multidigraph $G = (V, E)$ with *vertex* set V and *link* set E . Loops are not allowed, while more than one links can join the same pair of vertices. Each link $e \in E$ is associated with a *nonnegative, continuous, nondecreasing* latency function $\ell_e(\cdot)$, which specifies the time needed to traverse e as a function of the link congestion on e . We call G a *graph* with *edge* set E if it is undirected and a *digraph* with *arc* set E if it is directed.¹ Throughout the paper, by a path we mean a *simple* one, i.e., it contains *no repeated vertices*; by a path (or a cycle) in a digraph we mean a directed one. For any $u, v \in V$, a path in G from u to v is called a u - v path. For convenience, graphs and digraphs are collectively referred to as (di)graphs.

Let $k \geq 1$ be a positive integer. We write $[k]$ for the set $\{1, \dots, k\}$ of positive integers smaller than or equal to k . Given k distinct origin-destination pairs of vertices (s_i, t_i) in G , we call G a *k-commodity network* if for each $i \in [k]$, $s_i \neq t_i$ and G contains at least an s_i - t_i path. Given a *traffic demand* (vector) $\mathbf{r} = (r_i)_{i=1}^k$, the traffic in G comprises k flows, each for one commodity. For every $i \in [k]$, *the flow of commodity* i has an amount of r_i ; it is formed by an infinite number of players traveling from s_i to t_i . Each player selects a single path from his origin to his destination that has a minimum latency, given the congestion imposed by the rest of the players. Assuming a continuum of players, the choice of each individual player has a negligible impact on the experiences of others.

Formally, let the triple (G, \mathbf{r}, ℓ) denote a k -commodity selfish routing instance, where latency functions $\ell_e(\cdot)$, $e \in E$, are collectively represented by ℓ . The routing instance is sometimes written as $(G, (s_i, t_i)_{i=1}^k, \mathbf{r}, \ell)$ in order to explicitly specify the origin-destination pairs. For each $i \in [k]$, let \mathcal{P}_i be the set of s_i - t_i paths in G ; a *flow of*

¹Edges and arcs are collectively called links. An undirected link is an edge, and a directed link is an arc.

commodity i is a nonnegative vector $\mathbf{f}_i = (f_i(P))_{P \in \mathcal{P}_i}$ satisfying $\sum_{P \in \mathcal{P}_i} f_i(P) = r_i$. The combination of $\mathbf{f}_1, \dots, \mathbf{f}_k$ gives rise to a k -commodity flow $\mathbf{f} = (\mathbf{f}_i)_{i=1}^k$ for routing instance (G, \mathbf{r}, ℓ) . Under \mathbf{f} , each link e that is contained by some path in $\mathcal{P} = \cup_{i=1}^k \mathcal{P}_i$ experiences a congestion $f(e) = \sum_{i=1}^k \sum_{P \in \mathcal{P}_i: e \in P} f_i(P)$, and thus a link latency $\ell_e(f(e))$. Accordingly, each path P contained by $\cup_{Q \in \mathcal{P}} Q$ and any player traveling through P suffer from a path latency $\ell_P(\mathbf{f}) = \sum_{e \in P} \ell_e(f(e))$. In this nonatomic routing game, a Nash equilibrium is characterized by Wardrop’s principle [28] in a way that all players travel only on the minimum latency paths from their own origins to their own destinations.

Definition 1.1 A k -commodity flow \mathbf{f} in (G, \mathbf{r}, ℓ) is called a *Nash equilibrium (NE)* of (G, \mathbf{r}, ℓ) if for each $i \in [k]$ and each $P \in \mathcal{P}_i$ with $f_i(P) > 0$, it holds that $\ell_P(\mathbf{f}) = \min_{Q \in \mathcal{P}_i} \ell_Q(\mathbf{f})$.

By the classical result of Beckmann et al. [2] (see also [17, 23]), the NE of (G, \mathbf{r}, ℓ) exist, and are essentially unique in the sense that each link experiences the same latency under all NE of (G, \mathbf{r}, ℓ) . Thus, for each $i \in [k]$, the common latency experienced by all players traveling from s_i to t_i in any NE of (G, \mathbf{r}, ℓ) has the same value, which we denote by $\ell_i(G, \mathbf{r})$, and refer to as *the equilibrium latency of commodity i* for (G, \mathbf{r}, ℓ) .

1.2 Braess’s Paradox

The formal definition of Braess’s paradox involves the specific meaning of subnetworks. Let $(G, (s_i, t_i)_{i=1}^k, \mathbf{r}, \ell)$ be a k -commodity selfish routing instance. For any index set $\mathcal{K} \subseteq [k]$, a subgraph or subdigraph H of G is referred to as a *subnetwork* of (G, \mathcal{K}) if H is a $|\mathcal{K}|$ -commodity network with origin-destination pairs $(s_i, t_i)_{i \in \mathcal{K}}$. Following [17], a subnetwork of $(G, [k])$ is simply called a *subnetwork* of G .

Given routing instance (G, \mathbf{r}, ℓ) , it is natural to ignore the commodities of zero demands, and focus on $\mathcal{K}(\mathbf{r}) = \{i \in [k] : r_i > 0\}$. Let H be a subnetwork of $(G, \mathcal{K}(\mathbf{r}))$, and ℓ' be the restriction of ℓ to the link set of H . In a mild abuse of notation, we use (H, \mathbf{r}, ℓ) to denote the selfish routing instance $(H, (r_i)_{i \in \mathcal{K}(\mathbf{r})}, \ell')$, and, for each $i \in [k]$, use $\ell_i(H, \mathbf{r})$ to denote the equilibrium latency of commodity i for (H, \mathbf{r}, ℓ) . For simplicity, in case of single-commodity, i.e., $k = 1$, we often write \mathbf{r} as r (the traffic demand from the single origin to the single destination), and drop the subscript 1 in notions s_1, t_1, \mathcal{P}_1 , and $\ell_1(\cdot, r)$ in our discussions. When \mathbf{r} is an all-one vector we often write it as $\mathbf{1}$.

Definition 1.2 We say that *Braess’s paradox occurs* in (G, \mathbf{r}, ℓ) if there exists subnetwork H of $(G, \mathcal{K}(\mathbf{r}))$ and $h \in \mathcal{K}(\mathbf{r})$ such that $\ell_h(H, \mathbf{r}) < \ell_h(G, \mathbf{r})$ and $\ell_i(H, \mathbf{r}) \leq \ell_i(G, \mathbf{r})$ for all $i \in \mathcal{K}(\mathbf{r}) - \{h\}$. A k -commodity network G is said to be *Braess’s paradox free*, or *paradox-free* for short, if for any nonnegative traffic demand vector \mathbf{r} and any nonnegative, continuous, nondecreasing latency functions ℓ , Braess’s paradox does not occur in (G, \mathbf{r}, ℓ) . A k -commodity network that is not Braess’s paradox free is called *Braess’s paradox ridden*, or *paradox-ridden* for short.

We remark that the definition of paradox-ridden network here is a substantial relaxation of the ones given by Roughgarden [22] and Fotakis et al. [12], which admit instances suffering from *the most severe* performance loss in terms of Braess’s paradox.

For example, Braess’s paradox occurs in the routing instances (a) and (b) depicted in Fig. 1a and b, respectively, where one unit flow is to be routed from each origin to its corresponding destination. The latency function $\ell_e(x)$ on arc e of network G_a (a.k.a. directed Wheatstone network) or network G_b is either x or a constant 0, 1 or 2 as represented by the symbol beside the arc in the figure. In instance (a) (resp. (b)), the subnetwork H_a (resp. H_b) is obtained from G_a (resp. G_b) by deleting the dotted arc. It is easy to see that, in the single-commodity case (a), at the NE of $(G_a, 1, \ell)$ all players go through the path $suvt$, while at the NE of $(H_a, 1, \ell)$ half of players go through path sut and the other half go through path svt ; this gives $\ell(H_a, 1) = 1.5 < 2 = \ell(G_a, 1)$. Similarly, in the 2-commodity case (b), at the NE of $(G_b, \mathbf{1}, \ell)$ no player uses arc (s_1, t_1) (which has a constant latency 2 regardless of any congestion), while at the NE of $(H_b, \mathbf{1}, \ell)$ the flows of commodities 1 and 2 use disjoint paths s_1t_1 and s_2t_2 , respectively; it follows that $\ell_1(H_b, \mathbf{1}) = 2 = \ell_1(G_b, \mathbf{1})$ and $\ell_2(H_b, \mathbf{1}) = 1 < 2 = \ell_2(G_b, \mathbf{1})$. Hence, both network G_a and G_b are paradox-ridden. In contrast, the 3-commodity network G_c depicted in Fig. 1c is paradox-free, as our main result (Theorem 1.6) below guarantees.

A nice hereditary property implied by Definition 1.2 is that every paradox-free k -commodity network with origin-destination pairs $(s_i, t_i)_{i \in [k]}$ must be a paradox-free $|I|$ -commodity network with origin-destination pairs $(s_i, t_i)_{i \in I}$ for any nonempty $I \subseteq [k]$ (see Lemma 2.4). This property would have been lost if one modified the definition by requiring the existence of a subnetwork H of (G, \mathbf{r}) instead of $(G, \kappa(\mathbf{r}))$. Take the digraph G in Fig. 2 as an example. When G being considered as a 3-commodity network in (a), for any demand \mathbf{r} , every subnetwork of (G, \mathbf{r}) contains the path $s_1s_2t_2$, which would imply that G is paradox-free if one adopted the modified definition. However, G as a 2-commodity network in (b) is still paradox-ridden under the modified definition.

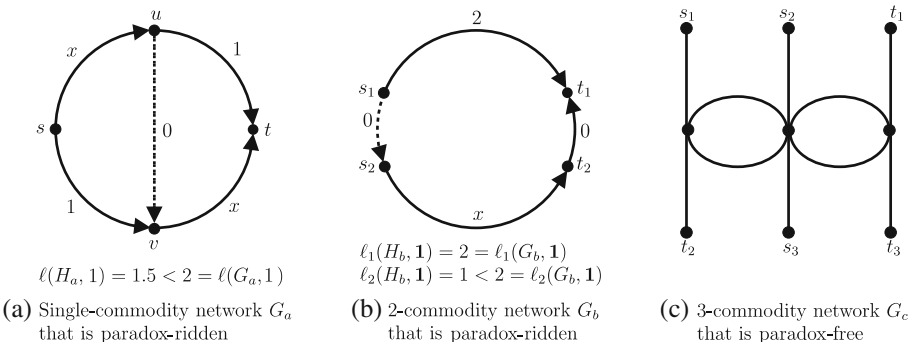


Fig. 1 Braess’s paradox in nonatomic selfish routing

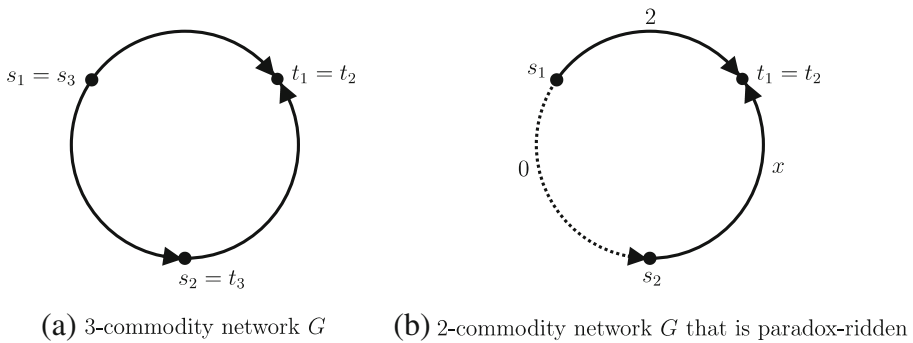


Fig. 2 The hereditary property of paradox-freeness

1.3 Our Contributions on Paradox-Free Networks

Considering a k -commodity network G with origin-destination pairs $(s_i, t_i)_{i=1}^k$, it would be convenient to think of G being *irredundant* in the sense that each link and each vertex of G are contained in at least an s_i - t_i path for some $i \in [k]$.

Milchtaich [20] established a series-parallel characterization for excluding Braess’s paradox in irredundant single-commodity undirected networks. The nice characterization confirms an unproven assertion of Murchland [21], and partially solved the open question of characterizing paradox-free networks, proposed by Roughgarden [22]. In this paper, we complete the solution by characterizing all irredundant k -commodity undirected and directed networks with the series-parallel and coincident conditions (see Theorem 1.6). In particular, our results on multicommodity networks answer the open question raised by Milchtaich [20]. The theoretical results imply polynomial time algorithms for recognizing paradox-free k -commodity undirected networks and paradox-free k -commodity planar directed networks.

Single-Commodity Networks The paradox-freeness of irredundant single-commodity networks is characterized in [20] for the undirected case and in this paper for the directed case (see Theorem 1.4 below), using the notion of two-terminal series-parallel networks.

Definition 1.3 An irredundant single-commodity network G with origin-destination pair (s, t) is said to be *two-terminal series-parallel*, or *s-t series-parallel* to be more specific, if one of the following conditions holds.

- (i) G is undirected, and there do not exist two s - t paths in G that pass a common edge in opposite directions.
- (ii) G is directed, and its underlying graph is s - t series-parallel.

Theorem 1.4 An irredundant single-commodity network is paradox-free if and only if it is two-terminal series-parallel.

Given the result by Milchtaich [20] for undirected networks, intuitively, one might expect a graphical characterization less restrictive than the two-terminal series-parallel one for directed networks, where players lose the flexibility to traverse a link in either direction, and Braess's paradox might have fewer chances to occur. However, the intuition is disproved by our result on directed networks.

Despite the similarity of the necessary and sufficient condition for graphs and that for digraphs in Theorem 1.4, the necessity proof for digraphs constitutes our first technical contribution. Milchtaich's model for single-commodity undirected network [20] allows a class of latency functions wider than ours (which is useful for proving the sufficiency). However the model does preclude predetermined directionality (which is a common feature of many real networks and modeled by digraphs [3]). Milchtaich's wider class of latency functions does not contribute the necessity proof of the characterization for directed networks. Neither the results nor their proofs in [20] can imply the necessity in the directed case. In the half-page necessity proof for graphs [20], given any undirected network that is not s - t series-parallel, Milchtaich derived a Braess's paradox by finding a special pair of vertices in the network whose existence relies on the property that the graph has two s - t paths which go through an edge in opposite directions. Such a property is lost when considering digraphs. This is the main hurdle to extending Milchtaich's proof to digraphs. New ideas and approaches are required to overcome the difficulty due to more complicated structures of digraphs.

- We use the existence of cycles in digraphs to obtain these special pairs of vertices (see Lemma 3.1), which in turn lead us to Braess's paradoxes (see Lemmas 3.4 and 3.3).
- We translate edge traverses in opposite directions in graphs into the existences of s - t paths through reversed arcs in digraphs. Instead of relying on special pairs of vertices as Milchtaich did, we derive Braess's paradox directly in acyclic digraphs that are not s - t series-parallel using an inductive argument that carefully exploits the properties of digraphs (see Theorem 3.5).

Multicommodity Networks To characterize paradox-free multicommodity network, we need the concept of block (see e.g., [26]). Let H be a (di)graph. We say that H is *2-connected* if it is connected and has no cut-vertices. A *block* of H is a maximal 2-connected sub(di)graph of H . If H is two-terminal series-parallel, then each block B of H is also two-terminal series-parallel, where the terminals of B are uniquely determined by the structure of H (see Lemma 2.8), and are referred to as *the terminals of B in H* .

For each $i \in [k]$, let $G_i = (V_i, E_i)$ be the maximum (in terms of the number of links) irredundant subnetwork of $(G, \{i\})$. Then G_i , consisting of all s_i - t_i paths in G , is an irredundant single-commodity network with origin-destination pair (s_i, t_i) . We call s_i and t_i the *terminals* of G_i .

Definition 1.5 If G_i and G_j (possibly $i = j$) are two-terminal series-parallel, B is a block of both G_i and G_j , and the set of terminals of B in G_i and that in G_j are the same, then B is called a *coincident block* of G_i and G_j .

In Definition 1.5, the orders of terminals (which is the origin and which is the destination) of B in G_i and G_j do not have to be the same. A common block B of G_i and G_j is coincident if G is directed, but it is not necessarily coincident if G is undirected, as in the undirected case the terminal sets of B in G_i and G_j might be different. In addition, the order of a coincident block’s terminals makes a difference between directed and undirected networks. Given a coincident block B of G_i and G_j , in the directed case, the order of B ’s terminals is unique; while in the undirected case, it is possible that the origin (resp. destination) of B in G_i is the destination (resp. origin) of B in G_j . An example of such different orders can be seen from G_1 and G_2 for $G = G_c$ in Fig. 1c.

The next theorem summarizes the main results of this paper, which concern both the directed and undirected case.

Theorem 1.6 (Main result) *An irredundant k -commodity network G is paradox-free if and only if G satisfies the following conditions:*

- (i) *Series-parallel Condition: for each $i \in [k]$, $G_i = (V_i, E_i)$ is an s_i - t_i series-parallel network; and*
- (ii) *Coincident Condition: for any $i, j \in [k]$, either $E_i \cap E_j = \emptyset$ or the (di)graph induced by $E_i \cap E_j$ consists of all coincident blocks of G_i and G_j .*

The theorem clearly generalizes Theorems 1.4. This generalization is by no means straightforward. Indeed, the coincident condition specifies the interactions of players with different origins and destinations, capturing in the context of paradox-freeness a key property of asymmetric nonatomic selfish routing, which, to the best of our knowledge, was not studied previously.

The configurations depicted in Figs. 1b and 3 give visualizations of how Braess’s paradox occurs in multicommodity directed networks and undirected networks, respectively, where $\mathfrak{u}(G_b)$ in Fig. 3a is the underlying graph of G_b in Fig. 1b.

Moreover, we show in Section 4 that the three configurations are typically all the forbidden structures for paradox-free multicommodity networks. This complements




Fig. 3 Braess’s paradoxes in 2-commodity instance (G, \mathbf{r}, ℓ) , where for $G \in \{\mathfrak{u}(G_b), G_d\}$ and $H = G \setminus e$ it holds that $\ell_1(H, \mathbf{1}) = 2 = \ell_1(G, \mathbf{1})$ and $\ell_2(H, \mathbf{1}) = 1 < 2 = \ell_2(G, \mathbf{1})$

the result that Wheatstone network (G_a in Fig. 1a and its underlying graph) is the only forbidden configuration for paradox-free single-commodity networks.

The proofs for the necessity of the coincident condition constitute our second technical contribution. In undirected networks, coping with the nonidentical sets of terminals for a common block of G_i and G_j turns out to be the key for obtaining the paradox configurations in Fig. 3.

- In connecting the different pairs of terminals, careful path selections to avoid unnecessary intersections, which involve techniques from graph connectivity theory, reveal the essence of the problem.

In directed networks, the difficulty lies on the fact that pure graph theory cannot enforce two intersecting blocks to be identical as it does for undirected networks.

- A more elaborate inductive method is applied to reduce the proof to smaller networks obtained by arc deletion and contraction. As arc contractions might create Braess's paradox which is not possessed by the original network, the challenging task is to guarantee that the selected arc contraction does not destroy the paradox-freeness (see the proof outline in Section 4.2 for more details).

Our results imply polynomial time algorithms for recognizing paradox-free k -commodity undirected networks and planar directed networks (see Corollary 5.2). In our proofs, the constructions of Braess's paradoxes only use linear latency functions of form $\ell_e(x) = a_e x + b_e$ for constants $a_e, b_e \geq 0$. An immediate corollary is that our results all hold even if the paradox-freeness is defined with respect to this kind of linear latency functions.

Our characterizations for paradox-free (paradox-ridden) networks show that the multicommodity cases are natural and nontrivial extensions from their single-commodity counterparts, which stands in contrast to the dichotomy between single- and 2-commodity networks in terms of severity of Braess's paradox [17, 22]. It is known that Braess's paradox can be dramatically more severe for multicommodity networks than single-commodity ones. More importantly, the characterization for paradox-free networks in the multicommodity case differs significantly from the single commodity case. Braess's paradox can occur in the multicommodity case even in simple series-parallel networks. An example is the 2-commodity network G_b in Fig. 1b that is paradox-ridden and series parallel.

1.4 Related Work

Most literature on characterizing network topologies for various properties of NE in selfish routing is restricted to the single-commodity case. The most related work is the aforementioned two-terminal series-parallel characterization by Milchtaich [20] for paradox-free undirected networks in nonatomic selfish routing. Recently, Cenciarelli et al. [5] obtained an equivalent characterization using the conception of "homeomorphic" from [20]. The authors proved that a single-commodity directed network is paradox-ridden if and only if it contains a subgraph homeomorphic to G_a in Fig. 1a. This work is independent from ours. The characterization with forbidden structure is valuable, and it is also verified by our results in Section 3. In [20], Milchtaich

proved that the networks which guarantee all NE to be weakly Pareto efficient are exactly those with *linearly independent routes*, meaning that every s - t path has at least an edge that does not belong to any other s - t path. The linear route independence was also shown to be a characterization for paradox-free single-commodity undirected networks with *heterogeneous* players, meaning that the latency functions are player-specific: the same link under the same congestion might give different latencies to different players using it. (Note that the linear independence result can not be extended to digraphs as shown by the paradox-ridden network G_a in Fig. 1a, where each of the three s - t paths has a unique arc that is not used by any other s - t path.) For the nonatomic routing with heterogeneous players, Milchtaich [18] characterized undirected networks such that, given any strictly increasing latency functions, each player's latency is the same at all NE. These networks are either nearly parallel or consist of two or more nearly parallel networks connected in series. In contrast to nonatomic routing, an *atomic routing* game has only a finite number of players, each controlling a noneligible portion of traffic. The most studied scenarios are *unsplittable routing* where each player routes his flow through a single path from his origin to his destination, and *routing with unit demands* where each player controls a unit of flow. Milchtaich [19] identified some sufficient conditions for directed network topologies that guarantee the existence of at least one pure NE in unsplittable atomic routing with unit demands for either player-specific latency functions or weighted players. The series connection of some small digraphs was shown to be sufficient for the existence of a pure NE.

As far as multicommodity networks are concerned, Holzman and Monderer [14] studied the unsplittable atomic routing game with unit demands that is played on a directed network with two distinguished vertices u and v , where every arc belongs to at least one u - v path. Note that such a constraint of networks is much stronger than ours of irredundant networks (defined at the beginning of Section 1.3); we only discard links never used by any player. Under this topological constraint, Holzman and Monderer [14] proved that the class of extension-parallel networks is exactly the one guaranteeing the existence of a strong equilibrium. This is a generalization of the previous result on the single-commodity case [15]. Epstein et al. [10] characterized the so-called *efficient* undirected networks in which all NE of each unsplittable routing instance with unit demands are socially optimal, i.e., they are among routings that minimize the maximum path latency between any origin-destination pair. It was shown that the efficient multicommodity networks are either trees or two vertices joined by parallel edges, while the efficient single-commodity networks are exactly those with linearly independent routes. The authors [10] also obtained characterizations of efficient undirected networks for the routing game where both individual players and the network (society) wish to minimize their own maximum edge latencies. Recently, Fujishige et al. [13] considered general nonatomic congestion games and gave a characterization of the maximal combinatorial property of strategy spaces for which Braess paradox does not occur. In a nutshell, bases of matroids are exactly this maximal structure.

Since its discovery [4] in 1968, Braess paradox has been a subject of considerable research on transportation networks [25], communication and computation systems [23], electrical circuits [6], and more general scientific context; see [22]

for an overview. In analyzing the worst-case severity of Braess's paradox in selfish routing games, Lin et al. [17] defined the Braess ratio of a k -commodity instance (G, \mathbf{r}, ℓ) as the maximum of $\min_{i=1}^k \ell_i(G, \mathbf{r})/\ell_i(H, \mathbf{r})$ over all subnetworks H of G . For nonatomic routing in single-commodity directed networks with $n \geq 2$ vertices, Roughgarden [22] established a tight upper bound $\lfloor n/2 \rfloor$ on the Braess ratio. The upper bound was generalized by Lin et al. [17] who showed that removing a set S of arcs from a single-commodity directed network G with origin-destination pair (s, t) can only decrease the equilibrium latency at most by a factor of $1 + c$, where c is the maximum number of arcs in S that form a matching of $G \setminus \{s, t\}$. In contrast to the linear upper bound for the single-commodity case, Lin et al. [17] proved an exponential lower bound on the Braess ratio for nonatomic routing in 2-commodity directed networks, and this bound is surprisingly attained by a one-arc removal.

Braess's paradox indicates the inefficiency of NE, which is usually measured by the *price of anarchy* (PoA), i.e., the worst-case ratio between the system objective value (of the maximum path latency between any origin-destination pair) at a NE and that of an optimal flow (which minimizes the maximum path latency). Strengthening the observation that the PoA is bounded below by the Braess ratio, Lin et al. [17] proved that the PoA of nonatomic routing is $n - 1$ for single-commodity directed network with n vertices, and $2^{O(\min\{kn, m \log n\})}$ for k -commodity directed network with n vertices and m arcs. A natural network design problem motivated by Braess's paradox is, given an nonatomic k -commodity routing instance (G, \mathbf{r}, ℓ) on directed network G with n vertices, to find a subnetwork H of $(G, \kappa(\mathbf{r}))$ such that $\max_{i \in \kappa(\mathbf{r})} \ell_i(H, \mathbf{r})$ is minimized. Assuming $P \neq NP$, the problem was shown to have very strong inapproximability. When $k = 1$ the best-possible approximation ratio for the problem is $\lfloor n/2 \rfloor$, implying the impossibility of detecting Braess's paradox (even in its worst manifestations) efficiently in directed single-commodity networks [22]. When $k \geq 2$ no polynomial time algorithm for the problem can achieve an approximation ratio $2^{o(n)}$, showing exponential inapproximability for detecting Braess's paradox in directed multicommodity networks [17]. In case of single commodity, the trivial algorithm that outputs the entire digraph G turns out to give the best possible approximation ratio (up to a constant factor) for the single-commodity case [22]; if, in addition, all latencies are linear, then the trivial algorithm guarantees a $(4/3)$ -approximation [22, 23]. Similar network design problems were later investigated for unsplittable atomic routing by Azar and Epstein [1] and for nonatomic single-commodity routing by Fotakis [12], where the efficiency was measured by the total latency [1] or the maximum link latency [12] of all players.

Organization The rest of the paper is organized as follows: In Section 2, we introduce notations and present preliminaries. In Sections 3 and 4, we characterize paradox-free single-commodity networks and paradox-free multicommodity networks, respectively. In Section 5, we give a wrap-up of the proof for our main result (Theorem 1.6), and then discuss its corollaries on recognition of paradox-free networks and the relations between the paradox-freeness of a network and that of its underlying graph or orientations. In Section 6, we conclude with remarks on directions of future research.

2 Preliminaries

We begin by introducing graph-theoretic notations used in the paper. Let $G = (V, E)$ be a (di)graph without loops that may possibly contain parallel links. As usual, G is said to be *acyclic* if G contains no cycle. If G is directed, we use $\mathfrak{u}(G)$ to denote the underlying (undirected) graph of G .

We often write $H \subseteq G$ to mean that H is a sub(di)graph of G . We use $V(H)$ and $E(H)$ to denote the vertex set and link set of H , respectively. For $u \in V(H) \cup E(H)$, we often write $u \in H$ if it is clear from the context that u is a vertex or it is a link.

Let $G_1, G_2 \subseteq G$. By $G_1 \cup G_2$, we mean the sub(di)graph of G with vertex set $V(G_1) \cup V(G_2)$ and link set $E(G_1) \cup E(G_2)$. If G_1 and G_2 have at least one vertex in common, by $G_1 \cap G_2$, we mean the sub(di)graph of G with vertex set $V(G_1) \cap V(G_2)$ and link set $E(G_1) \cap E(G_2)$. If $V(G_1) \cap V(G_2) = \emptyset$, we often use short notation $G_1 \cap G_2 = \emptyset$. If $G_1 \subseteq G_2$ and $G_2 \subseteq G_1$, we write $G_1 = G_2$.

Let $U \subseteq V \cup E$. The sub(di)graph of G obtained from G by removing elements in U is denoted by $G \setminus U$. In case of U consisting of a singleton u , we write $G \setminus u$ instead of $G \setminus \{u\}$. If G is undirected, we write an edge of G with end-vertices u, v as uv . If G is directed, we write an arc of G with tail u and head v as (u, v) .

Let $P = v_1 v_2 \dots v_h$ be a v_1 - v_h path in G .² For any $1 \leq i < j \leq h$, we say that v_i *precedes* v_j in P , or equivalently v_j *follows* v_i in P ; such a relation is written as $v_i \prec_P v_j$. If $1 \leq i \leq j \leq h$, we write $v_i \leq_P v_j$, and use $P[v_i, v_j]$ to denote the subpath of P from v_i to v_j . For convenience, we set $P(v_i, v_j) = P[v_i, v_j] \setminus v_i$, $P[v_i, v_j) = P[v_i, v_j] \setminus v_j$ and $P(v_i, v_j) = P[v_i, v_j] \setminus \{v_i, v_j\}$. Let $u, v \in V$, and Q be a u - v path in G . We say that P and Q are internally vertex-disjoint if $V(P) \cap V(Q) \subseteq \{v_1, v_h\} \cap \{u, v\}$.

Lemma 2.1 *Let $G = (V, E)$ be an acyclic digraph. If $u, v, w \in V$, and P and Q are u - v path and v - w path in G , respectively, then $P \cup Q$ is a u - w path in G .*

Consider graph G . For any $X, Y, Z \subseteq V$, we call a path in G between a vertex in X and a vertex in Y an X - Y path. We say that vertices in Z *separate* X from Y if $G \setminus Z$ does not contains any X - Y path.

Theorem 2.2 (Menger's theorem) *If $G = (V, E)$ is a graph and $X, Y \subseteq V$, then the minimum number of vertices that separate X from Y in G is equal to the maximum number of vertex-disjoint X - Y paths in G .*

Corollary 2.3 *Let $G = (V, E)$ be a 2-connected graph with $u, v, w, x, y \in V$. If $u \neq v, w \neq x$ and $wx \in E$, then (i) G contains a u - v path going through wx , and (ii) G contains a u - v path going through y .*

Proof Since G is 2-connected, no single vertex in V can separate $\{u, v\}$ from $\{w, x\}$. It follows from Theorem 2.2 that there are two vertex-disjoint $\{u, v\}$ - $\{w, x\}$ paths P

²We emphasize again that all paths in this paper are simple. They are all acyclic.

and Q in G . Note that $P \cup wx \cup Q$ is a u - v path containing wx . Thus (i) holds. It is clear that (ii) follows from (i). \square

Lemma 2.4 *Let G be a k -commodity network (directed or undirected) with origin-destination pairs $(s_i, t_i)_{i=1}^k$. If G is paradox-free, then for any nonempty $I \subseteq [k]$, each $|I|$ -commodity network $G' (\subseteq G)$ with origin-destination pairs $(s_i, t_i)_{i \in I}$ is paradox-free.*

Proof Suppose on the contrary that there exist $I \subseteq [k]$ and $|I|$ -commodity network G' with origin-destination pairs $(s_i, t_i)_{i \in I}$ such that Braess's paradox occurs in (G', \mathbf{r}', ℓ') for some $\mathbf{r}' = (r'_i)_{i \in I}$ and ℓ' . We construct an instance (G, \mathbf{r}, ℓ) by defining \mathbf{r} and ℓ as follows:

- $r_i = r'_i$ for all $i \in I$, and $r_i = 0$ for all $i \in [k] - I$;
- $\ell_e(x) = \ell'_e(x)$ for each $e \in E(G')$, and $\ell_e(\cdot) = \infty$ for each $e \in E(G) - E(G')$.

It is easy to see that $\kappa(\mathbf{r}) = \kappa(\mathbf{r}')$, (G, \mathbf{r}, ℓ) and (G', \mathbf{r}', ℓ') have the same NE flow (in a way of ignoring the commodities with zero demands), and $\ell_i(G, \mathbf{r}) = \ell'_i(G', \mathbf{r}')$ for all $i \in \kappa(\mathbf{r})$. As Braess's paradox occurs in (G', \mathbf{r}', ℓ') , there exist $H \subseteq G' \subseteq G$ and $h \in \kappa(\mathbf{r}') = \kappa(\mathbf{r})$ such that $\ell_h(H, \mathbf{r}) = \ell'_h(H, \mathbf{r}') < \ell'_h(G', \mathbf{r}') = \ell_h(G, \mathbf{r})$ and $\ell_i(H, \mathbf{r}) = \ell'_i(H, \mathbf{r}') \leq \ell'_i(G', \mathbf{r}') = \ell_i(G, \mathbf{r})$ for all $i \in \kappa(\mathbf{r}') - \{h\} = \kappa(\mathbf{r}) - \{h\}$, showing that Braess's paradox occurs in (G, \mathbf{r}, ℓ) . The contradiction to paradox-freeness of G proves the lemma. \square

The remainder of this section is devoted to the discussion on two-terminal series-parallel (di)graphs. The way of defining the series-parallel property [20] in terms of link traverse directions (see Definition 1.3(i)) is valid only for (undirected) graphs. It can not be extended to digraphs. On the other hand, the series-parallel property can be uniformly defined for both graphs and digraphs in the following recursive way.

Definition 2.5 A single-commodity network G with origin-destination pair (s, t) is *two-terminal series-parallel* or *s - t series-parallel* if

- (i) G has a single link with ends s and t ; or
- (ii) G is obtained by connecting two smaller o_i - d_i series-parallel networks H_i , $i = 1, 2$, *in series* – identifying d_1 and o_2 , and naming o_1 as s , and d_2 as t ; or
- (iii) G is obtained by connecting two smaller o_i - d_i series-parallel networks H_i , $i = 1, 2$, *in parallel* – identifying o_1 and o_2 to form s and identifying d_1 and d_2 to form t .

In the above definition, the vertices s and t are called the *terminals* of the series-parallel (di)graph G . The equivalence of Definitions 2.5 and Definition 1.3 was proved by Milchtaich [20].³ The above recursive definition implies polynomial time

³Milchtaich's proof [20] concerned only undirected graphs. The directed case is simply an immediate corollary.

recognition algorithms for recognizing two-terminal series-parallel (di)graphs [27]. It is well-known that each series-parallel graph of at least two edges has either a degree 2 vertex or a pair of parallel edges. The following observation on two-terminal series-parallel digraphs will be useful in our proofs.

Lemma 2.6 *Let G be a two-terminal series-parallel digraph with at least two arcs. Then G contains either a pair of parallel arcs or a vertex whose indegree and outdegree are both one.*

Proof We prove the lemma by induction on the number m of arcs in G . The base case where G has exactly two arcs is trivial. We assume that $m \geq 3$ and the lemma holds for all two-terminal series-parallel networks with at least 2 and at most $m - 1$ arcs. Since $m \geq 2$, there exist two smaller two-terminal series-parallel digraphs H_1 and H_2 whose connection in series or in parallel gives G . Suppose without loss of generality that H_1 has at least two arcs. It follows that H_1 satisfies our lemma, because H_1 has fewer edge than G . Clearly, any pair of parallel arcs of H_1 is the one in G . Moreover, if a vertex v has exactly one in-neighbor and exactly one out-neighbor in H_1 , then v cannot be a terminal of H_1 , and therefore cannot have any neighbors in H_2 . Thus G satisfies the lemma. \square

Corollary 2.7 *Let $G = (V, E)$ be an irredundant single-commodity undirected network with origin-destination pair (s, t) . For any distinct vertexes $u, v \in V$, G contains two vertex-disjoint $\{s, t\}$ - $\{u, v\}$ paths.*

Proof Since $u \neq v$, and each of u and v is contained in an s - t path of G , it is easy to see that no single vertex of G can sperate $\{s, t\}$ from $\{u, v\}$. The conclusion is instant from Theorem 2.2. \square

For any s - t series-parallel (di)graph G , let $U(G)$ denote the set of its cut-vertices, and $\mathcal{B}(G)$ the set of its blocks. Let s - t block graph of G be defined as a graph with vertex set $\{s, t\} \cup U(G) \cup \mathcal{B}(G)$ and edge set $\{vB : v \in \mathcal{B}, v \in \{s, t\} \cup U(G), B \in \mathcal{B}(G)\}$. Note that the definition of s - t block graphs here is slightly different from that of standard block graphs in textbook (see, e.g., page 56 of [9]). The next observation concerns with the block-chain structure of two-terminal series-parallel (di)graphs. See Fig. 4 for an illustration of the directed case. (The undirected case is visualized by ignoring the arc directions.)




Fig. 4 The block structures of two-terminal series-parallel networks

Lemma 2.8 *Let G is an s - t series-parallel (di)graph. If $p = |\mathcal{B}(G)|$, then the s - t block graph of G is a path of form $v_1 B_1 v_2 \dots v_i B_i v_{i+1} \dots v_p B_p v_{p+1}$ such that the following hold:*

- (i) $v_1 = s$, $v_{p+1} = t$, and $U(G) = \{v_2, \dots, v_p\}$ if and only if $U(G) \neq \emptyset$ if and only if $p \geq 2$;
- (ii) each s - t path in G goes through v_1, v_2, \dots, v_p in this order;
- (iii) $\mathcal{B}(G) = \{B_1, \dots, B_p\}$;
- (iv) for each $i \in [p]$, B_i is a v_i - v_{i+1} series-parallel (di)graph; and
- (v) for each $i \in [p]$ and each v_i - v_{i+1} path P , either B_i is a link or B_i contains a v_i - v_{i+1} path Q such that P and Q are internally disjoint.

Proof By Definition 2.5, it suffices to consider G obtained by connecting two smaller o_i - d_i series-parallel networks H_i ($i = 1, 2$) in series or in parallel. In the former case, the result follows from applying inductions on H_1 and H_2 . In the latter case, G is 2-connected, $\mathcal{B}(G)$ is G itself, and the statements (i)–(iv) are trivially true. To see (v), suppose without loss of generality that $P \subseteq H_1$. Then we may take Q to be any o_2 - d_2 path, i.e., s - t path, in H_2 . \square

3 Single-Commodity Networks

In view of Milchtaich's characterization for undirected networks [20], in this section we mainly focus on single-commodity directed networks, while leaving a brief discussion on the undirected case to the end of this section.

We study single-commodity directed network $G = (V, E)$ with origin-destination pair (s, t) . Without loss of generality, we may assume that G is *irredundant*, which particularly implies that G is connected. We will characterize paradox-freeness of G with s - t series-parallel property. Our efforts are mainly devoted to proving the necessity of the property.

In Milchtaich's proof for undirected networks (Proposition 1 of [20]), the embedding of $\mathfrak{u}(G_a)$ (see Fig. 1a for G_a) in a single-commodity undirected network implies the network is paradox-ridden, while such an embedding is implied by a special pair of distinct vertices u, v such that u precedes v in an s - t path and follows v in another s - t path of the network. To derive such u, v , Milchtaich used the property that the graph has two s - t paths which go through an edge in opposite directions. Such a property is lost when considering digraphs. This is the main hurdle to extending Milchtaich's proof to digraphs.

To get around the difficulty, our first step is to use the existence of a cycle to derive that of a special pair of vertices in G (Lemma 3.1), which in turn gives the occurrence of Braess's paradox (Lemmas 3.3 and 3.4). Our second step is to derive Braess's paradox in acyclic G that is not s - t series-parallel. We translate edge traverses in opposite directions in graphs into the existences of s - t paths through reversed arcs in digraphs; instead of relying on the special pair of vertices, we derive Braess's paradox directly using an inductive argument that carefully exploits the properties of digraphs (Theorem 3.5).

Lemma 3.1 *If there is a cycle in G , then there exist s - t paths $P, Q \in \mathcal{P}$ and distinct vertices $u, v \in V(P) \cap V(Q)$ such that u precedes v in P and u follows v in Q , i.e., $u \prec_P v$ and $v \prec_Q u$.*

Proof Assume on the contrary that no such P, Q and u, v stated in the lemma exist. Let $C = v_1 v_2 \dots v_m v_1$ be a cycle in G . For each $2 \leq i \leq m$, let $C[v_1, v_i]$ stand for the unique v_1 - v_i path in the directed cycle C . Since G is an irredundant network, every arc (v_i, v_{i+1}) of C is contained in some s - t path $Q_i \in \mathcal{P}$, $i = 1, 2, \dots, m$, where $(v_m, v_{m+1}) = (v_m, v_1)$.

If $Q_1[s, v_2)$ and $Q_2(v_2, t]$ have some vertex in common, say u , then our contradiction assumption is violated by $u \prec_{Q_1} v_2$ and $v_2 \prec_{Q_2} u$. Hence $Q_1[s, v_2]$ and $Q_2[v_2, t]$ intersect only at vertex v_2 , and their concatenation forms an s - t path $Q_1[s, v_2] \cup Q_2[v_2, t]$ containing $C[v_1, v_3]$, which we write as $Q_{1,2}$. Then we consider $Q_{1,2}[s, v_3]$ and $Q_3[v_3, t]$. They, by our contradiction assumption, intersect only at v_3 (otherwise $u \prec_{Q_{1,2}} v_3$ and $v_3 \prec_{Q_3} u$ for some vertex $u \in Q_{1,2}[s, v_3] \cap Q_3(v_3, t]$), and their concatenation forms an s - t path $Q_{1,2,3} := Q_{1,2}[s, v_3] \cup Q_3[v_3, t]$ containing $C[v_1, v_4]$. Continuing this way, we concatenate $Q_{1,2,\dots,i}[s, v_{i+1}]$ and $Q_{i+1}[v_{i+1}, t]$ to obtain an s - t path $Q_{1,2,\dots,i+1}$ containing $C[v_1, v_{i+2}]$ for $i = 1, 2, \dots, m - 2$ inductively. In particular $Q_{1,2,\dots,m-1}$ contains $C[v_1, v_m]$, implying that v_1 precedes v_m in $Q_{1,2,\dots,m-1} \in \mathcal{P}$. However v_1 follows v_m in $Q_m \in \mathcal{P}$. The contradiction to our assumption establishes the lemma. \square

Definition 3.2 We call G an s - t paradox if $G = P_1 \cup P_2 \cup P_3$ is the union of three paths P_1, P_2 and P_3 with the following properties:

- (i) P_1 is an s - t path going through distinct vertices a, u, v, b in this order, i.e., $s \preceq_{P_1} a \prec_{P_1} u \prec_{P_1} v \prec_{P_1} b \preceq_{P_1} t$;
- (ii) P_2 is an a - v path with $V(P_2) \cap V(P_1) = \{a, v\}$;
- (iii) P_3 is a u - b path with $V(P_3) \cap V(P_1) = \{u, b\}$ and $V(P_3) \cap V(P_2) = \emptyset$.

The s - t paradox defined above is often denoted as (P_1, P_2, P_3) . As illustrated in Fig. 5, P_2 and P_3 are vertex-disjoint, and they intersect with P_1 only at their ends. Observe that G_a in Fig. 1a is an s - t paradox $(suv t, sv, ut)$ with $s = a$ and $t = b$.

Lemma 3.3 *If G contains an s - t paradox, then G is paradox-ridden.*




Fig. 5 An s - t paradox (P_1, P_2, P_3)

Proof Let $G' = (V', E') = P_1 \cup P_2 \cup P_3$ be an s - t paradox (P_1, P_2, P_3) contained in $G = (V, E)$. In the subdigraph G' , let e_1 and e_2 be the two outgoing arcs from a with $e_1 \in P_1$ and $e_2 \in P_2$, and let e_3 and e_4 be the two incoming arcs to b with $e_3 \in P_1$ and $e_4 \in P_3$. We define routing instance $(G, 1, \ell)$ by $\ell_e(x) = x$ if $e \in \{e_1, e_3\}$, $\ell_e(x) = 1$ if $e \in \{e_2, e_4\}$, $\ell_e(x) = 0$ if $e \in E' - \{e_1, e_2, e_3, e_4\}$, and $\ell_e(x) = \infty$ if $e \in E - E'$. See Fig. 5 for an illustration. The unique NE in $(G, 1, \ell)$ sends the one-unit flow all through path P_1 and suffers from a latency $\ell(G, 1) = \ell_{e_1}(1) + \ell_{e_3}(1) = 2$. Let subnetwork H of G be obtained from G by removing an arc e on $P_1[u, v]$. Then the unique NE in $(H, 1, \ell)$ splits the one-unit flow equally at vertex a , sending a half via path $P_2[a, v] \cup P_1[v, t]$ and the other half via path $P_1[a, u] \cup P_3[u, b] \cup P_1[b, t]$. This incurs a latency $\ell(H, 1) = 1 + 0.5 = 1.5$, showing a Braess's paradox. \square

Lemma 3.4 *If there is a cycle in G , then G contains an s - t paradox.*

Proof By Lemma 3.1, there exist s - t paths $P, Q \in \mathcal{P}$ and distinct vertices $u, v \in V(P) \cap V(Q)$ such that u precedes v in P and u follows v in Q . Let such P, Q, u, v be taken such that $Q[v, u]$ is as long as possible. It follows from the maximality of $Q[v, u]$ that

(i) $Q[s, v]$ and $P[u, t]$ are vertex-disjoint.

Suppose otherwise $Q[s, v]$ and $P[u, t]$ have a common vertex w . Note from $u \in Q(v, t]$ that $w \neq u$. It follows that we should have chosen P, Q, u, w instead of P, Q, u, v because $Q[w, u]$ properly contains $Q[v, u]$. So (i) holds. A symmetric argument shows the following,

(ii) $Q(u, t]$ and $P[s, v]$ are vertex-disjoint.

Consider a traverse of Q from s to t . Let a be the last vertex on $Q[s, v]$ with $a \in P$, meaning $Q(a, v) \cap P = \emptyset$; let b be the first vertex on $Q(u, t]$ with $b \in P$, meaning $Q(u, b) \cap P = \emptyset$. It follows from (i) that $a \in P[s, u)$, and from (ii) that $b \in P(v, t]$ (See Fig. 5 for an illustration with $P_1 = P, P_2 = Q[a, v]$ and $P_3 = Q[u, b]$). It can be seen from the choices and positions of a and b that $P, Q(a, v)$ and $Q(u, b)$ are pairwise vertex-disjoint, implying that $P \cup Q[a, v] \cup Q[u, b]$ is an s - t paradox $(P, Q[a, v], Q[u, b])$ contained in G . \square

For any arc subset $K \subseteq E$, let $G\langle K \rangle$ be the digraph obtained from G by reversing the directions of all arcs in K . The set of reversed arcs is written as \bar{K} .

Theorem 3.5 *If G does not contain any s - t paradox, then G is s - t series-parallel.*

Proof As G contains no s - t paradox, it follows from Lemma 3.4 that G is acyclic. Therefore by Lemma 2.1, we have

(i) for any $u, v, w \in V$ and any u - v path P_{uv} and v - w path P_{vw} in G , $P_{uv} \cup P_{vw}$ is a u - w path in G .

By Definition 1.3, digraph G is s - t series-parallel if and only if its underlying graph $\mathfrak{u}(G)$ has no pair of s - t paths which go through some edge in opposite

directions. Recall that each arc of G is contained in some s - t path in G . To justify the theorem, it suffices to prove the equivalent statement that for any nonempty subset K of E , there is no s - t path in $G\langle K \rangle$ going through some arc in \bar{K} .

By contradiction, we assume on the contrary that there is $\emptyset \neq K \subseteq E$ with minimum $|K|$ such that $G\langle K \rangle$ contains an s - t path P going through some arc in \bar{K} . The minimality of K implies that $\bar{K} \subseteq P$, and \bar{K} induces a number of, say m , subpaths of P from v_i to w_i , $i = 1, 2, \dots, m$, where $v_i <_P w_i <_P v_{i+1} <_P w_{i+1}$ for all $i = 1, 2, \dots, m - 1$. Notice that

- (ii) the reverse of $P[v_i, w_i]$ is a path $\bar{P}[w_i, v_i] = (V_i, E_i)$ in G for each $i \in [m]$;
- (iii) K is the disjoint union of E_1, E_2, \dots, E_m ;
- (iv) $P[w_i, v_{i+1}]$, $i = 0, 1, \dots, m$ are paths in G , where $w_0 = s$ and $v_{m+1} = t$;
- (v) for any $i \in [m]$, there exists a v_i - v_1 path Q_i in G .

Statements (ii)–(iv) are straightforward observations. See Fig. 6 for an illustration. We prove (v) by induction on i . The base case $i = 1$ is trivial. Suppose that $2 \leq i \leq m$ and there is a v_h - v_1 path Q_h in G for each $h \in [i - 1]$.

Since G is an irredundant network, it contains a v_i - t path Q . If Q and $P[s, v_i]$ are vertex-disjoint, then $P[s, v_i] \cup Q$ is an s - t path in $G(\cup_{h=1}^{i-1} E_h)$ going through $\cup_{h=1}^{i-1} \bar{E}_h$. However $\emptyset \neq \cup_{h=1}^{i-1} E_h \subsetneq \cup_{h=1}^m E_h$ and (iii) imply a contradiction to the minimality of K . Hence

- $P[s, v_i]$ and $Q(v_i, t)$ have some vertex in common, say c .

Observe from (iv) and (i) that $P[w_{i-1}, v_i] \cup Q$ is a path in G , implying that $P[w_{i-1}, v_i]$ and Q are vertex-disjoint, and

- $c \in P[s, w_{i-1}] \cap Q(v_i, t)$.

Note that $c \in P[w_{h-1}, w_h]$ for some $h \leq i - 1$. By (iv) and (ii), the subdi-graph $P[w_{h-1}, v_h] \cup \bar{P}[w_h, v_h]$ of G contains a c - v_h path, written as R . Consider the concatenation of R and the v_h - v_1 path Q_h (whose existence is guaranteed by the inductive hypothesis). We see from (i) that $R \cup Q_h$ is a c - v_1 path in G . In turn we concatenate the v_i - c path $Q[v_i, c]$ and the c - v_1 path $R \cup Q_h$ into a v_i - v_1 path in G , which establishes (v).

Applying (ii) and (v) with $i = m$, we deduce from (i) that $S_{w_m v_1} = \bar{P}[w_m, v_m] \cup Q_m$ is a w_m - v_1 path in G . Since G is an irredundant network, it contains an s - w_m path $S_{s w_m}$ and a v_1 - t path $S_{v_1 t}$. By (i), $S = S_{s w_m} \cup S_{w_m v_1} \cup S_{v_1 t}$ is an s - t path in G satisfying $s <_S w_m <_S v_1 <_S t$. Observe that $w_m, v_1 \in P \cap S$. This enables us to take distinct vertices $u, v \in P \cap S$ such that

- $u <_S v$ and $v <_P u$,
- $P[v_1, w_m] \subseteq P[v, u]$, and
- $P[v, u]$ is as long as possible.




Fig. 6 The path P in $G\langle K \rangle$, where arcs in G and K are drawn as solid and dotted ones, respectively

Notice from $s \leq_S u <_S v \leq_S t$ that $v \neq s$ and $u \neq t$. Recall from (iv) that $P[s, v_1] \cup P[w_m, t] = P[w_0, v_1] \cup P[w_m, v_{m+1}] \subseteq G$. From $P[v_1, w_m] \subseteq P[v, u]$, we derive

$$(vi) \quad P[s, v] \cup P[u, t] \subseteq G.$$

The maximality of $P[v, u]$ enforces that

$$(vii) \quad P[s, v] \cap S[u, t] = \emptyset \text{ and } P(u, t) \cap S[s, v] = \emptyset.$$

Otherwise we should have taken v to be some vertex $v' \in P[s, v] \cap S(u, t)$, or u to be some vertex $u' \in P(u, t) \cap S[s, v]$. See Fig. 6 for the positions of the contradictory v' and u' on P .

Consider a traverse of P from s to t . Let a be the last vertex of $P[s, v]$ with $a \in S$, meaning that $P(a, v) \cap S = \emptyset$ (such a vertex a exists because $s \in P[s, v] \cap S$ as $v \neq s$). Let b be the first vertex of $P(u, t]$ with $b \in S$, meaning $P(u, b) \cap S = \emptyset$ (such a vertex b exists because $t \in P(u, t] \cap S$ as $u \neq t$). It follows from (vii) that $a \in S[s, u)$ and $b \in S(v, t]$, giving $a <_S u <_S v <_S b$. Thus, by (vi), we see that $(S, P[a, v], P[u, b])$ is an s - t paradox in G . The contradiction completes the proof of the theorem. \square

We are now ready to prove the directed case of Theorem 1.4, restated as follows.

Theorem 3.6 *An irredundant single-commodity directed network G is paradox-free if and only if it is two-terminal series-parallel.*

Proof If G is paradox-free, then, by Lemma 3.3, G contains no s - t paradox, which along with Theorem 3.5 implies that G is s - t series-parallel.

Conversely, if G is s - t series-parallel, then $\mathfrak{u}(G)$ is s - t series-parallel by Definition 1.3. It can be seen from the definition that $\{\mathfrak{u}(P) : P \in \mathcal{P}\}$ is exactly the set of s - t paths in $\mathfrak{u}(G)$. Suppose for a contradiction that there exist $H \subseteq G$, traffic demand r and nonnegative, continuous, nondecreasing latency functions $\ell_e, e \in E$ such that $\ell(H, r) < \ell(G, r)$. For each edge e of $\mathfrak{u}(G)$, set $\underline{\ell}_e(\cdot) = \ell_{e'}(\cdot)$ where e is the underlying edge of the arc $e' \in E$. Then $\underline{\ell}(\mathfrak{u}(H), r) = \ell(H, r) < \ell(G, r) = \underline{\ell}(\mathfrak{u}(G), r)$, exhibiting a Braess’s paradox in $\mathfrak{u}(G)$. We deduce from the undirected case of Theorem 1.4 that $\mathfrak{u}(G)$ is not s - t series-parallel. The contradiction proves the theorem. \square

The sufficiency stated in the above theorem has an alternative proof based on Milchtaich’s work [20]. Milchtaich’s model for single-commodity undirected network allows a class of latency functions wider than ours; in particular, two (different) functions are defined for each edge e , one specifying the latency of passing e from its end u to its end v and the other specifying the latency of passing e from v to u . This class of latency functions brings a kind of directionality into Milchtaich’s model – Regardless of how the edges in a two-terminal series-parallel undirected network are directed and their latencies are defined, Braess’s paradox does not occur. To be more specific, given any routing instance in a single-commodity directed network G , let us consider its corresponding instance in the underlying graph $\mathfrak{u}(G)$ under Milchtaich’s

model. For each edge of $\mathfrak{u}(G)$, the latency function of its corresponding arc is associated with passing the edge in the direction the same as the arc, and an infinite latency function is given to passing the edge in the opposite direction. Then Braess's paradox occurs in the routing instance in $\mathfrak{u}(G)$ if and only if so does in the routing instance in G . If G is two-terminal series-parallel, then so is $\mathfrak{u}(G)$. Milchtaich's result [20] implies that the undirected instance does not admit any Braess's paradox, and we derive the same for the directed instance.

A Remark on Single-Commodity Undirected Networks As mentioned above, the class of latency functions investigated in Milchtaich's paper [20] is wider than the one in this paper. Milchtaich defined Braess's paradox free undirected networks with respect to this wider class of functions, and characterized these networks with the two-terminal series-parallel property. While the sufficiency of Theorem 1.4 (the undirected case) is immediate from Milchtaich's original theorem (Theorem 1 of [20]), the necessity can only follow from Milchtaich's proof, where only our class of latency functions (i.e., nonnegative, continuous, and nondecreasing ones) was used for constructing Braess's paradox.

4 Multicommodity Networks

In this section we study k -commodity network $G = (V, E)$ with origin-destination pairs $(s_i, t_i)_{i=1}^k$, where $k \geq 2$. For convenience, we assume that G is an irredundant network. We are to prove that the series-parallel condition and coincident condition (see Theorem 1.6) are necessary and sufficient for G to be paradox-free.

By virtue of the block chain structure of series-parallel (di)graphs, the sufficiency proof, which builds on the result for single-commodity (Theorem 1.4), turns out to be easier than the necessity part.

Theorem 4.1 *Let G be an irredundant multicommodity network. If G satisfies the series-parallel condition and the coincident condition, then G is paradox-free.*

Proof Given any nonnegative traffic demand vector $\mathbf{r} = (r_i)_{i=1}^k$ and nonnegative, continuous and nondecreasing latency functions ℓ , suppose that H is a subnetwork of $(G, \kappa(\mathbf{r}))$. Let $\mathbf{g} = (\mathbf{g}_i)_{i=1}^k$ and $\mathbf{h} = (\mathbf{h}_i)_{i=1}^k$ be the NE of (G, \mathbf{r}, ℓ) and (H, \mathbf{r}, ℓ) , respectively.

Consider an arbitrary commodity $i \in \kappa(\mathbf{r})$. Since G_i is s_i - t_i series-parallel, by Lemma 2.8, we suppose that the s_i - t_i block graph of G_i is $v_1 B_1 v_2 \dots v_p B_p v_{p+1}$, where $s_i = v_1$ and $t_i = v_{p+1}$.

For each $q \in [p]$, recall from Lemma 2.8(iv) that B_q is a v_q - v_{q+1} series-parallel (di)graph. We consider it as a single-commodity network with origin-destination pair (v_q, v_{q+1}) . Set traffic demand $r^{(q)} = \sum_{j: B_q \subseteq G_j} r_j$. Let $\ell^{(q)}$ denote the restriction of ℓ to the links in B_q . Since H is a subnetwork of $(G, \kappa(\mathbf{r}))$, it contains at least an s_i - t_i path. Every such s_i - t_i path intersects B_q with a v_q - v_{q+1} path, implying that $H_q = B_q \cap H$ is nonempty. Furthermore, it follows from the coincidence condition that H_q contains all v_q - v_{q+1} paths in H . Since \mathbf{g} is the NE for

(G, \mathbf{r}, ℓ) , its restriction to B_q , denoted as $\mathbf{g}^{(q)}$, is the NE for $(B_q, r^{(q)}, \ell^{(q)})$, where $g^{(q)}(P) = \sum_{j:P \subseteq G_j} \sum_{Q:P \subseteq Q \in \mathcal{P}_j} g_j(Q)$ for each v_q - v_{q+1} path in B_q . Similarly, the restriction of \mathbf{h} to H_q , denoted as $\mathbf{h}^{(q)}$, is the NE for $(H_q, r^{(q)}, \ell^{(q)})$, where $h^{(q)}(P) = \sum_{j:P \subseteq G_j} \sum_{Q:P \subseteq Q \in \mathcal{P}_j, Q \subseteq H} h_j(Q)$ for each v_q - v_{q+1} path in H_q . As B_q is v_q - v_{q+1} series-parallel, it follows from Theorem 1.4 that B_q is paradox-free, giving $\ell_P^{(q)}(\mathbf{g}^{(q)}) = \ell_i(B_q, r^{(q)}) \leq \ell_i(H_q, r^{(q)}) = \ell_Q^{(q)}(\mathbf{h}^{(q)})$ for any v_q - v_{q+1} paths P, Q with $g^{(q)}(P) > 0$ and $h^{(q)}(Q) > 0$.

Now we have $\ell_i(G, \mathbf{r}) = \sum_{q \in [p]} \ell^{(q)}(B_q, r^{(q)}) \leq \sum_{q \in [p]} \ell^{(q)}(H_q, r^{(q)}) = \ell_i(H, \mathbf{r})$. By the arbitrary choice of $i \in \kappa(\mathbf{r})$, we deduce that G is paradox-free. \square

For the necessity proof, we have the following observation, which provides us with series-parallel structure of each individual $G_i, i \in [k]$ for further study on the coincident condition.

Corollary 4.2 *If G is paradox-free, then G satisfies the series-parallel condition.*

Proof It is straightforward from Lemma 2.4 that all $G_i (i \in [k])$ are paradox-free, and then by Theorem 1.4, they all are two-terminal series-parallel. \square

The rest of this section is devoted to the proof of paradox-freeness implying the coincident condition. In view of Lemma 2.4, we may focus on 2-commodity networks: the undirected ones are investigated in Section 4.1, while the directed ones are discussed in Section 4.2. In both undirected and directed cases, we derive Braess’s paradox assuming the coincident condition fails.

4.1 The Undirected Case

In this subsection, we investigate *irredundant* 2-commodity undirected network G . The following Lemma 4.3 plays an important role in our proof. It helps to verify the fact that the intersecting blocks of G_1 and G_2 are their common blocks. Then we cope with the different terminal sets of G_1 and G_2 by connecting them via carefully selected paths, and use these paths to construct Braess’s paradox as desired.

Lemma 4.3 *For any $i \in [2]$ and any distinct vertices $u, v \in V_i$, all u - v paths of G are contained in G_i .*

Proof By contradiction, let u - v path P be a counterexample to the lemma with a minimum number of edges. Then all vertices of $P(u, v)$ are outside G_i . Recall from Corollary 2.7 that G_i contains vertex-disjoint $\{s_i, t_i\}$ - $\{u, v\}$ paths Q_1 and Q_2 with $u \in Q_1$ and $v \in Q_2$. It follows that $Q_1 \cup P \cup Q_2$ is an s_i - t_i path in G . Thus $Q_1 \cup P \cup Q_2 \subseteq G_i$ by the definition of G_i . This gives a contradiction to the fact that P and G_i are arc-disjoint. \square

Theorem 4.4 *Let G be an irredundant 2-commodity undirected network. If G is paradox-free, then it satisfies the coincident condition.*

Proof Notice from Corollary 4.2 that G_i is an s_i - t_i series-parallel network for $i = 1, 2$. Suppose that $E_1 \cap E_2 \neq \emptyset$. Then we have an edge $uv \in E_1 \cap E_2$, a block B_1 of G_1 with terminals a, b and a block B_2 of G_2 with terminals c, d such that $uv \in B_1 \cap B_2$. Since B_1 is 2-connected, by Corollary 2.3(i), each edge of B_1 is on some u - v path of B_1 . Thus Lemma 4.3 enforces that $B_1 \subseteq G_2$. As B_1 is 2-connected and have common edge uv with block B_2 of G_2 , we have $B_1 \subseteq B_2$. A symmetric argument gives $B_2 \subseteq B_1$, implying $B_1 = B_2$. It remains to prove that $\{a, b\} = \{c, d\}$. Assume the contrary that $\{a, b\} \neq \{c, d\}$, which particularly implies that $B_1 = B_2$ is not an edge. We distinguish between two cases depending on whether there exist an a - b path and a c - d path with one containing the other. We derive contradictions in both cases by constructing Braess’s paradoxes.

Case 1 There are an a - b path and a c - d path in B_1 such that one contains the other. Suppose without loss of generality that B_1 contains an a - b path P with $c, d \in P$. As B_1 is not an edge ab , it follows from Lemma 2.8(v) that there is an a - b path Q in B_1 which is internally disjoint from P , i.e., $P \cup Q$ is a cycle.

Take an arbitrary edge $e_1 \in P[c, d]$ and an arbitrary edge $e_2 \in Q$. Let latency functions ℓ be defined as follows: $\ell_{e_1}(x) = x, \ell_{e_2}(x) = 2, \ell_e(x) = 3$ for all edges e inside B_1 but outside $P \cup Q$, and $\ell_e(x) = 0$ for edges e outside B_1 and all edges e inside $(P \cup Q) \setminus \{e_1, e_2\}$. See Fig. 7a for an illustration. Define \mathbf{r} by $r_1 = r_2 = 1$.

On one hand, in the NE for (G, \mathbf{r}, ℓ) , when restricted to $B_1 = B_2$, commodity 1 only uses path P and commodity 2 only uses path $P[c, d]$. It follows that $\ell_1(G, \mathbf{r}) = \ell_2(G, \mathbf{r}) = 2$. On the other hand, consider the subgraph H obtained from G by removing an edge e_3 inside P but outside $P[c, d]$ (such an edge does exist because $\{a, b\} \neq \{c, d\}$). Since e_3 belongs to the cycle $P \cup Q \subseteq G$, it can be seen that $H = G \setminus e_3$ is a subnetwork of G . Observe that in any NE for (H, r, ℓ) , when restricted to $B_1 \cap H = B_2 \cap H$, commodity 1 only uses path Q and commodity 2 only uses path $P[c, d]$. It follows that $\ell_1(G, \mathbf{r}) = 2$ and $\ell_2(G, \mathbf{r}) = 1$, showing the occurrence of Braess’s paradox in (G, \mathbf{r}, ℓ) .

Case 2 No a - b path in B_1 contains any c - d path and no c - d path in B_1 contains any a - b path. It follows that no a - b path in B_1 contains $\{c, d\}$ and no c - d path in B_1 contains $\{a, b\}$. If $a = d$, then by Corollary 2.3(ii), there is an a - b path in B_1 that passes through c , and this a - b path clearly contains a c - d path, a contradiction to the hypothesis of Case 2. Therefore $a \neq d$, and symmetric arguments provide $\{a, b\} \cap \{c, d\} = \emptyset$. Moreover, by Corollary 2.3(ii), we may take P to be an a - b path in B_1 containing c and Q to be an a - b path in B_1 containing d .




Fig. 7 Braess’s paradoxes in 2-commodity undirected networks

Consider a traverse of P from a to b . Let u denote the last common vertex of $P[a, c]$ and Q , and v denote the first common vertex of $P[c, b]$ and Q . Obviously $u \neq v$, $c \in P[u, v]$, and $P[u, v] \cap Q$ contains only two vertices u and v . Swapping u and v if necessary, we obtain

- (i) There exist distinct vertices $u, v \in P \cap Q$ such that $a \preceq_Q u \preceq_Q v \preceq_Q b$, $c \in P(u, v)$ and $P[u, v] \cap Q$ contains only u and v

See Fig. 7b for an illustration. Since no a - b path in B_1 contains $\{c, d\}$, statement (i) enforces

- (ii) $d \in Q(u, v)$.

By (i) and (ii), we may take edge $e_1 \in Q[u, d]$ and edge $e_2 \in P[c, v]$. Let latency functions ℓ be defined as follows: $\ell_{e_1}(x) = x$, $\ell_{e_2}(x) = 2$, $\ell_e(x) = 3$ for all edges e inside B_1 but outside $P[u, v] \cup Q$, and $\ell_e(x) = 0$ for edges e outside B_1 and all edges e inside $(P[u, v] \cup Q) \setminus \{e_1, e_2\}$. Define \mathbf{r} by $r_1 = r_2 = 1$. Then in the NE for (G, \mathbf{r}, ℓ) , when restricted to $B_1 = B_2$, commodity 1 only uses the a - b path Q , and commodity 2 only uses the c - d path $P[c, u] \cup Q[u, d]$. It follows that $\ell_1(G, \mathbf{r}) = \ell_2(G, \mathbf{r}) = 2$. On the other hand, consider the subgraph H obtained from G by removing an edge $e_3 \in Q[d, v]$ (such an edge does exist because of (ii)). It can be seen that H is a subnetwork of G , and in the NE for (H, \mathbf{r}, ℓ) , when restricted to $B_1 \cap H = B_2 \cap H$, commodity 1 only uses path $Q[a, u] \cup P[u, v] \cup Q(v, b)$ and commodity 2 only uses path $P[c, u] \cup Q[u, d]$. It follows that $\ell_1(G, \mathbf{r}) = 2$ and $\ell_2(G, \mathbf{r}) = 1$, showing the occurrence of Braess’s paradox in (G, \mathbf{r}, ℓ) . □

Remark 4.5 Suppose that G_1 and G_2 are both two-terminal series-parallel networks, and $E_1 \cap E_2$ induces all coincident blocks of G_1 and G_2 . It is easy to prove that these coincident blocks must be consecutive in the sequence of blocks of G_1 (resp. G_2) ordered according to the s_1 - t_1 block graph of G_1 (resp. the s_2 - t_2 block graph of G_2).

4.2 The Directed Case

In this subsection, we prove that paradox-freeness implies the coincident condition for *irredundant* 2-commodity directed network G . If G_1 and G_2 are two-terminal series-parallel, then, due to the directionality of G , coincident blocks of G_1 and G_2 are exactly their common blocks, because each block of G_i ($i = 1, 2$) has a unique source, which must be its origin, and a unique sink, which must be its destination. In the proof of Theorem 4.4, Lemma 4.3 helps to obtain a quick proof for $B_1 = B_2$. As the lemma does not extend to digraphs, more elaborative operations of arc deletion and contraction will be employed in directed networks for reaching the conclusion of $B_1 = B_2$ (i.e., its failure implies occurrence of Baress’s paradox).

Given a 2-commodity network D (not necessarily irredundant), we use D_i ($i = 1, 2$) to denote the maximum irredundant subnetwork of $(D, \{i\})$. The *maximum irredundant subnetwork* of D is $D_1 \cup D_2$. Instead of focusing on irredundant directed networks with origin-destination pairs $(s_i, t_i)_{i=1,2}$, our proof considers the set \mathcal{C} of

2-commodity directed networks D (not necessarily irredundant) whose maximum irredundant subnetworks $D_1 \cup D_2$ satisfy the coincident condition: if a block B_1 of D_1 and a block B_2 of D_2 have some arc in common, then B_1 and B_2 are *identical*, i.e., $B_1 = B_2$. Our goal is to prove $G \in \mathcal{C}$ provided G is paradox-free.

Our argument employs arc deletion and contraction frequently. Let D be a digraph and e an arc of D . We write $D \setminus e$ and D/e for the graphs obtained from D by deleting e and contracting e , respectively. If D' is a subdigraph of $D \setminus e$, then both notations $D' \setminus e$ and D'/e mean D' itself.

Proof Outline To prove $G \in \mathcal{C}$, we use the minimum counterexample approach. Assuming $G \notin \mathcal{C}$ is a minimum (in term of the number of arcs) counterexample, for any $e \in E$, we have $G \setminus e$ and G/e belong to \mathcal{C} whenever they are paradox-free subnetworks of G .

- The basic idea is to derive from $G \setminus e \in \mathcal{C}$ or $G/e \in \mathcal{C}$ either $G \in \mathcal{C}$ or a Braess's paradox in G , contradicting the assumption that $G \notin \mathcal{C}$ and G is paradox-free.

However, not all $e \in E$ can help us to fulfill this task, because G/e might have many s_i - t_i paths P ($i = 1, 2$) such that neither P nor $P \cup \{e\}$ is an s_i - t_i path in G , meaning that G/e might be paradox-ridden and we could not have $G/e \in \mathcal{C}$. This is where the inductive series-parallel property helps; it gives us a pair of arcs e_1, e_2 which are in parallel or in series in G_1 (recalling Lemma 2.6). The case of parallel e_1, e_2 is easy. Most of our efforts are devoted to investigating series e_1, e_2 . The argument is divided into the following main steps to prove:

- G_2 contains e_1 or e_2 , say e_1 .
- e_1 is the only incoming or outgoing arc at the common end v of e_1 and e_2 .
- G_1 and G_2 have a common arc other than e_1 (otherwise a Braess's paradox is obtained).
- for $i = 1, 2$, the set of s_i - t_i path in G_i/e_1 naturally correspond the set of s_i - t_i paths in G_i , implying G/e_1 is paradox-free and thus $G/e_1 \in \mathcal{C}$.

Then the final contradiction $G \in \mathcal{C}$ is reached due to $G/e_1 \in \mathcal{C}$ and $e_1 \in G_1 \cap G_2$.

Theorem 4.6 *Let G be a 2-commodity directed network. If G is paradox-free, then $G \in \mathcal{C}$.*

Proof By contradiction, let $G = (V, E) \notin \mathcal{C}$ be a counterexample to the theorem with minimum $|E|$. (The minimality of G instantly implies that G is irredundant.) Hence there exist a block B_i of $G_i = (V_i, E_i)$, $i = 1, 2$ such that B_1 and B_2 have some arc e^* in common, but $B_1 \neq B_2$. It can be assumed without loss of generality that

$$B_1 \not\subseteq B_2,$$

which implies $|E(B_1)| \geq 2$ because $e^* \in B_1 \cap B_2$. Since G is paradox-free, it follows from Corollary 4.2 that both G_1 and G_2 are two-terminal series-parallel. In turn, Lemma 2.8(iv) says that both B_1 and B_2 are two-terminal series-parallel. We break our proof into a sequence of claims.

Claim 1 $B_1 \not\subseteq G_2$.

Otherwise, B_1 is contained by some block of G_2 . This block must be B_2 because $e^* \in B_1$ and B_2 is the (unique) block of G_2 that contains e^* . However $B_1 \subseteq B_2$ gives a contradiction to $B_1 \not\subseteq B_2$,

Claim 2 $G \setminus e$ is paradox-free for any $e \in E$.

If there exists $e \in E$ such that Braess's paradox occurs in $(G \setminus e, \mathbf{r}, \ell)$ for some \mathbf{r} and ℓ . Extending ℓ by $\ell_e(x) = \infty$, Braess's paradox occurs in (G, \mathbf{r}, ℓ) , a contradiction to the assumption that G is paradox-free.

Claim 3 There exist distinct vertices u, v, w , and arcs $e_1 = (u, v)$, $e_2 = (v, w)$ in B_1 such that e_1 is the only incoming arc at v in G_1 , and e_2 is the only outgoing arc from v in G_1 .

If it were not the case, then Lemma 2.6 applies to the two-terminal series-parallel network B_1 with at least two arcs, and enforces that B_1 contains a pair of parallel arcs e_1 and e_2 . Switching e_1 and e_2 if necessary, we may assume $e_1 \neq e^*$. Clearly, $D = G \setminus e_1$ is a 2-commodity network with origin-destination pairs $(s_i, t_i)_{i=1,2}$. Since D is paradox-free by Claim 2, we deduce from the minimality of $G \notin \mathcal{C}$ that $D \in \mathcal{C}$. Therefore $B_1 \setminus e_1$ and $B_2 \setminus e_1$ are identical because they are blocks of D_1 and D_2 , respectively, and both contain e^* . Notice from $e_2 \in B_1 \setminus e_1 = B_2 \setminus e_1$ that $e_1 \in B_2$. Now $e_1 \in B_1 \cap B_2$ combined with $B_1 \setminus e_1 = B_2 \setminus e_1$ implies $B_1 = B_2$. The contradiction to $B_1 \not\subseteq B_2$ justifies Claim 3.

Claim 4 B_1/e_i is 2-connected for $i = 1, 2$.

The claim is an immediate corollary of Claim 3 and the 2-connectivity of B_1 .

Claim 5 G_2 contains either e_1 or e_2 .

Assuming the contrary, let D be obtained from $G \setminus \{e_1, e_2\}$ by adding a new arc $e_0 = (u, w)$. It is clear that D is still a 2-commodity network with origin-destination pairs $(s_i, t_i)_{i=1,2}$, $D_1 = G_1 \setminus \{e_1, e_2\} \cup \{e_0\}$, $B_1 \setminus \{e_1, e_2\} \cup \{e_0\}$ is the block of D_1 that contains e^* (note that $e^* \notin \{e_1, e_2\}$) and $B_2 \subseteq D_2$.

If Braess's paradox occurs in some routing instance (D, \mathbf{r}, ℓ') , then Braess's paradox occurs in (G, \mathbf{r}, ℓ) , where $\ell_{e_i}(x) = \ell'_{e_0}(x)/2$ for $i = 1, 2$ and $\ell_e(x) = \ell'_e(x)$ for all $e \in E - \{e_1, e_2\}$. The paradox-freeness of G implies that D is paradox-free, and the minimality of $G \notin \mathcal{C}$ implies $D \in \mathcal{C}$. It follows from $e^* \in B_2 \subseteq D_2$ that $B_1 \setminus \{e_1, e_2\} \cup \{e_0\} \subseteq D_2$.

As $e_0 \in D_2$, there is an s_2 - t_2 path P in D_2 with $e_0 \in P$. If $v \in P[s_2, u]$ or $v \in P[w, t_2]$, then, correspondingly, $P[s_2, v] \cup (vw) \cup P[w, t_2]$ or $P[s_2, u] \cup (uv) \cup P[v, t_2]$ is an s_2 - t_2 path in G containing e_2 or e_1 , showing a contradiction to our

assumption that $e_1, e_2 \notin G_2$. So $v \notin P$ and $P \setminus e_0 \cup \{e_1, e_2\} = P[s_2, u] \cup (uvw) \cup P[w, t_2]$ is an s_2 - t_2 path in G , again a contradiction to $e_1, e_2 \notin G_2$. Thus Claim 5 holds.

Reversing the directions of all arcs if necessary, we may assume without loss of generality that $e_1 \in G_2$. Since v is incident with both an incoming arc $e_1 \in E_1 \cap E_2$ and an outgoing arc $e_2 \in E_1$, we see that

$$v \notin \{s_1, t_1, s_2\}. \tag{1}$$

Claim 6 e_1 is the only incoming arc at v in G .

Assume on the contrary that there exists an arc $e_0 \in E - \{e_1\}$ which is incoming at v . Since $G = G_1 \cup G_2$, it follows from Claim 3 that $e_0 \in G_2$. Note that $e_0 \notin G_1$. As $e_1 \in G_2$, we consider an s_2 - t_2 path in G_2 that passes through e_1 . This path clearly avoids e_0 ; it is thus an s_2 - t_2 path in $D = G \setminus e_0$. It follows that D is a 2-commodity network with origin-destination pairs $(s_i, t_i)_{i=1,2}$, and $e_1 \in D_2$. Since D is paradox-free by Claim 2, and it has fewer arcs than G , we have $D \in \mathcal{C}$. Note that $D_1 = G_1$, $D_2 \subseteq G_2 \setminus e_0$, and B_1 is the block of D_1 containing e_1 . We deduce from $D \in \mathcal{C}$ and $e_1 \in D_2$ that B_1 is the block of D_2 which contains e_1 . It follows that $B_1 \subseteq D_2 \subseteq G_2 \setminus e_0 \subseteq G_2$, a contradiction to Claim 1. Hence Claim 6 holds.

Claim 7 $E(B_1) \cap E(G_2) \neq \{e_1\}$.

Suppose on the contrary that $E(B_1) \cap E(G_2) = \{e_1\}$. We construct routing instance (G, \mathbf{r}, ℓ) in which Braess’s paradox occurs. Recall that B_1 is a two-terminal series-parallel network with at least two arcs. Let a and b be the origin and destination of B_1 , respectively. It follows from Lemma 2.8(v) that B_1 contains two internally-disjoint a - b paths P and Q such that $e_1 \in P$. By Claim 3, we have $e_2 \in P$. Take e_3 to be an arbitrary arc in Q . Let R be an s_2 - t_2 path in G_2 that goes through e_1 . Let latency functions ℓ be defined as follows (see Fig. 8):

- $\ell_{e_1}(x) = x$,
- $\ell_{e_3}(x) = 2$,
- $\ell_e(x) = 0$ for all arcs $e \in P \cup Q \cup R \setminus \{e_1, e_3\}$,
- $\ell_e(x) = \infty$ for all arcs $e \in B_1 \setminus E(P \cup Q)$, and
- $\ell_e(x) = 2$ for all arcs $e \in G \setminus E(R \cup B_1)$.




Fig. 8 A Braess’s paradox in 2-commodity directed network

Note that the latencies of arcs outside B_1 are constants, either 0 or 2. Thinking of them as arc lengths, let α and β denote the length of a shortest s_1 - a path in $G_1 \setminus E(B_1)$ and that of a shortest b - t_1 path in $G_1 \setminus E(B_1)$, respectively.

Let the demand vector be defined as $\mathbf{r} = (1, 1)$. In the NE of (G, \mathbf{r}, ℓ) , commodity 1 chooses path P in B_1 and chooses an s_1 - a path of latency α and a b - t_1 path of latency β outside B_1 ; commodity 2 chooses path R . We have $\ell_1(G, \mathbf{r}) = \alpha + 2 + \beta$ and $\ell_2(G, \mathbf{r}) = 2$. Now consider the subdigraph $H = G \setminus e_2$. In the NE of (H, \mathbf{r}, ℓ) , commodity 1 chooses Q in B_1 , giving $\ell_1(H, \mathbf{r}) = \alpha + 2 + \beta$; commodity 2 chooses R , giving $\ell_2(H, \mathbf{r}) = 1$. This shows that Braess’s paradox occurs in (G, \mathbf{r}, ℓ) . The contradiction to the condition that G is paradox-free verifies Claim 7.

Instantly, $e_1 \in G_2$ and Claims 6 and 7 imply that either $u \neq s_2$ or $v \neq t_2$. It follows that $D = G/e_1$ is a 2-commodity network with origin-destination pairs $(s_i, t_i)_{i=1,2}$. For $i = 1, 2$, let \mathcal{Q}_i denote the set of s_i - t_i paths in D . By abuse of notation, we shall identify a digraph with its arc set. Claim 6 enables us to construct for $i = 1, 2$ a 1-1 correspondence φ_i between \mathcal{Q}_i and \mathcal{P}_i as follows: For any $Q \in \mathcal{Q}_i$.

$$\varphi_i(Q) = \begin{cases} Q & \text{if } Q \in \mathcal{P}_i; \\ Q \cup \{e_1\} & \text{otherwise.} \end{cases} \tag{2}$$

Claim 8 For $i = 1, 2$, φ_i is a 1-1 correspondence between \mathcal{Q}_i and \mathcal{P}_i .

To see φ_i is a mapping from \mathcal{Q}_i and \mathcal{P}_i , it suffices to consider the case where $Q \notin \mathcal{P}_i$ for some $Q \in \mathcal{Q}_i$. From $D = G/e_1$, we deduce that Q (as a subset of E) consists of two vertex disjoint paths $Q[s_i, u']$ and $Q[v', t_i]$ in G , where $\{u', v'\} = \{u, v\}$. If $u' = v$, then $v \notin \{s_1, s_2\}$ in (1) implies that $v = u'$ has an incoming arc on $Q[s_i, u']$ which is different from e_1 , a contradiction to Claim 6. So it must be the case that $u' = u$ and $v' = v$, giving $Q \cup \{e_1\} \in \mathcal{P}_i$.

It is clear that φ_i is injective. To see it is also surjective, we consider an arbitrary $P \in \mathcal{P}_i$. In case of $e_1 \notin P$, Claim 6 and (1) imply $v \notin P$. Therefore $P \in \mathcal{Q}_i$, and $\varphi_i(P) = P$. In case of $e_1 \in P$, we have $P/e_1 \in \mathcal{Q}_i$, and $\varphi_i(P/e_1) = P$. This proves Claim 8.

For $i = 1, 2$, the above proof also gives the inverse of φ_i as follows: For any $P \in \mathcal{P}_i$.

$$\varphi_i^{-1}(P) = \begin{cases} P & \text{if } e_1 \notin P; \\ P/e_1 & \text{otherwise.} \end{cases} \tag{3}$$

Claim 9 $D = G/e_1$ is paradox-free.

Suppose on the contrary that Braess’s paradox occurs in (D, \mathbf{r}, ℓ') for some \mathbf{r} and ℓ' , i.e., there exists a subnetwork F of $(D, \kappa(\mathbf{r}))$ and $\{g, h\} = \{1, 2\}$ such that

$$\ell'_g(F, \mathbf{r}) < \ell'_g(D, \mathbf{r}) \text{ and } \ell'_h(F, \mathbf{r}) \leq \ell'_h(D, \mathbf{r}). \tag{4}$$

Define $\ell_{e_1}(x) = 0$ and $\ell_e(x) = \ell'_e(x)$ for every $e \in E - \{e_1\}$. It is clear that \mathbf{r} is not a zero vector. We show that Braess’s paradox occurs in (G, \mathbf{r}, ℓ) .

For each subnetwork X of $(D, \kappa(r))$, we write $\varphi(X)$ for the subdigraph $\cup_{i=1,2}(\cup_{Q \subseteq X, Q \in \mathcal{Q}_i} \varphi_i(Q))$ of G (recalling Claim 8). Since G is an irredundant network, we have $\varphi(X) = G$ if and only if $X = D$. Given any flow $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2)$ for (X, \mathbf{r}, ℓ') , let $\varphi(\mathbf{x})$ be the flow $\mathbf{y} = (\mathbf{y}_1, \mathbf{y}_2)$ for $(\varphi(X), \mathbf{r}, \ell)$ defined by: $y_i(\varphi(Q)) = x_i(Q), i = 1, 2$, for each $Q \in \mathcal{Q}_i$ with $Q \subseteq X$. It follows from the definition of ℓ that $\ell_{\varphi(Q)}(\mathbf{y}) = \ell'_Q(\mathbf{x})$ for all $Q \in \mathcal{Q}_i (i = 1, 2)$ with $Q \subseteq X$. Hence, it can be seen from Definition 1.1 that $\varphi(\mathbf{x})$ is a NE for $(\varphi(X), \mathbf{r}, \ell)$ whenever \mathbf{x} is a NE for (X, \mathbf{r}, ℓ') .

Considering the NE flows \mathbf{d} and \mathbf{f} for (D, \mathbf{r}, ℓ') and (F, \mathbf{r}, ℓ') , respectively, we see that $\varphi(\mathbf{d})$ is the NE for $(\varphi(D), \mathbf{r}, \ell) = (G, \mathbf{r}, \ell)$ with $\ell_i(G, \mathbf{r}) = \ell'_i(D, \mathbf{r})$ for $i = 1, 2$, and $\varphi(\mathbf{f})$ is the NE for $(\varphi(F), \mathbf{r}, \ell)$ with $\ell_i(\varphi(F), \mathbf{r}) = \ell'_i(F, \mathbf{r})$ for $i = 1, 2$. It follows from (4) that Braess’s paradox occurs in (G, \mathbf{r}, ℓ) . The contradiction establishes Claim 9.

Claim 10 $D_i = G_i/e_1$ for $i = 1, 2$.

On one hand, for each arc e of D_i , there exists $Q \in \mathcal{Q}_i$ such that $e \in Q$. It is clear from (2) that $e \in \varphi_i(Q) \in \mathcal{P}_i$ and $e \in \varphi_i(Q)/e_1 \subseteq G_i/e_1$. On the other hand, for each arc e of G_i/e_1 , there exists $P \in \mathcal{P}_i$ such that $e \in P$. Therefore $e \in \varphi^{-1}(P) \in \mathcal{Q}_i$ by (3), saying $e \in D_i$. So Claim 10 holds.

The combination of Claim 9 and the minimality of $G \notin \mathcal{C}$ gives $D = G/e_1 \in \mathcal{C}$. Recall from Claim 7 that B_1/e_1 and G_2/e_1 have some arc in common. By Claim 10 it is equivalent to saying that B_1/e_1 and D_2 have some arc in common.

Observe that B_1/e_1 is a subdigraph of $G_1/e_1 = D_1$ (by Claim 10). Recall from Claim 4 that B_1/e_1 is 2-connected. So B_1/e_1 is contained in a block B of D_1 , which have some common arc with D_2 . It follows from $D \in \mathcal{C}$ that B is a block of $D_2 = G_2/e_1$. Now $B_1/e_1 \subseteq G_2/e_1$ along with $e_1 \in G_2$ implies $B_1 \subseteq G_2$. The contradiction to Claim 1 completes the proof of Theorem 4.6. □

Remark 4.7 Suppose that G_1 and G_2 are both two-terminal series-parallel directed networks and $E_1 \cap E_2$ induces all common blocks of G_1 and G_2 . Then the connected components of $G_1 \cap G_2$ can be ordered as C_1, \dots, C_h for some $h \geq 1$ such that each s_1 - t_1 path in G_1 passes these components in order of C_1, \dots, C_h , and each s_2 - t_2 path in G_2 passes these components in order of C_h, \dots, C_1 . (See Fig. 9 for an illustration).




Fig. 9 An instance for Remark 4.7

5 The Unified Characterization and Corollaries

The paradox-freeness of all irredundant networks (directed, undirected, single-commodity, multicommodity) has a unified characterization, as specified in Theorem 1.6. We first complete the proof of the characterization, then discuss its corollaries.

5.1 The Wrap-Up Proof for Irredundant Networks

We now summarize the proofs for the sufficiency and necessities of series-parallel condition (i) and coincidence condition (ii) in Theorem 1.6.

Proof of Theorem 1.6 In view of Theorem 1.4, we assume $k \geq 2$. The sufficiency of conditions (i) and (ii) and the necessity of condition (i) have been established in Theorem 4.1 and Corollary 4.2, respectively. To see the necessity of condition (ii), suppose that G is paradox-free. By Lemma 2.4, for any distinct $i, j \in [k]$, $G_i \cup G_j$ is a paradox-free irredundant 2-commodity network. In turn Theorem 4.4 or Theorem 4.6 applies to $G_i \cup G_j$, validating condition (ii). \square

5.2 Polynomial Time Recognition for Paradox-Freeness

We start this subsection with the discussion on single-commodity network $G = (V, E)$ with origin-destination pair (s, t) . Without loss of generality, we may assume that G is connected. Note that Milchtaich characterized only irredundant networks [20], and the result (the undirected case of Theorem 1.4) is easily extended to all undirected single-commodity networks as follows.

Theorem 5.1 *A single-commodity undirected network G is paradox-free if and only if G can be decomposed into a number of graphs $G'_i = (V'_i, E'_i)$, $i = 0, \dots, p$ for some $p \geq 0$ such that*

- (i) E is the disjoint union of E'_0, \dots, E'_p ;
- (ii) $|V'_i \cap V'_j| = 1$ for any $1 \leq i \leq p$, and $V'_i \cap V'_j = \emptyset$ for any $1 \leq i < j \leq p$;
- (iii) G'_0 is an irredundant single-commodity network with origin-destination pair (s, t) ; and
- (iv) G'_0 is s - t series-parallel.

Proof Suppose that G admits a decomposition into G'_0, \dots, G'_p such that (i)–(iv) hold. It is easy to see that no s - t path in G goes through any edge in $\cup_{i=1}^p E'_i$. So G is paradox-free if and only if G'_0 is. Indeed, by Theorem 1.4, it follows from (iii) and (iv) that G'_0 is paradox-free, proving the sufficiency.

Suppose that G is paradox-free. Let $G'_0 = (V'_0, E'_0) \subseteq G$ be the maximum (in terms of the number of edges) irredundant subnetwork of G . If $E'_0 = E$ then let $p = 0$, else let $G'_i = (V'_i, E'_i)$, $i = 1, \dots, p$ be the connected components in $G \setminus E'_0$ each of which contains at least one edge. It is obvious that (i) and (iii) hold, and the validity of (iv) is guaranteed by Theorem 1.4 because G is paradox-free. It remains to verify (ii). Notice from the definitions of G'_1, \dots, G'_p that $V'_i \cap V'_j = \emptyset$ for any

$1 \leq i < j \leq p$. If $|V'_0 \cap V'_i| \geq 2$ for some $1 \leq i \leq p$, then there exist distinct vertices $u, v \in V'_0 \cap V'_i$ and u - v path P in G'_i such that P and G'_0 have only u, v in common. In view of (iii) and (iv), the application of Corollary 2.7 to G'_0 provides two vertex-disjoint $\{s, t\}$ - $\{u, v\}$ paths Q and R in G'_0 . From $V(P) \cap V'_0 = \{u, v\}$ we deduce that $P \cup Q \cup R$ is an s - t path in G going through P . The maximality of G'_0 implies $P \subseteq G'_0$. However, $u \neq v$ says that P contains a least an edge, which belongs to E'_i ; now $E'_i \cap E'_0 = \emptyset$ gives a contradiction. So (ii) also holds, proving the necessity. \square

Recall that two-terminal series-parallel networks can be recognized in polynomial time [27]. Theorem 5.1 provides a polynomial time algorithm for the problem of determining whether a given undirected single-commodity network is paradox-free. Note that the single vertex, denoted as v_i , of $V'_0 \cap V'_i$ for any $2 \leq i \leq p$ in Theorem 5.1 is a cut-vertex of G , and $G'_i \setminus v_i$ consists of a number of components of $G \setminus v_i$ containing neither s nor t .

Algorithm 1 Recognizing Paradox-free single-commodity undirected networks

Input: undirected single-commodity network G with origin-destination pair (s, t) .

Output: the maximum irredundant subnetwork G' of G , and the assertion of G being paradox-free (“Yes”) or not (“No”)

1. $G' \leftarrow G, \quad S \leftarrow$ the set of cut-vertices of G
 2. **for** each $v \in S$ **do**
 3. remove from G' all vertices in the components of $G \setminus v$ containing neither s nor t (if any)
 4. **end-for**
 5. **if** G' is s - t series-parallel **then** output G' and “Yes” **else** output G' and “No”
-

Now, let G be a k -commodity network with origin-destination pairs $(s_i, t_i)_{i=1}^k$, and G_i the maximum irredundant subnetwork of $(G, \{i\})$, $i \in [k]$. The first corollary of Theorem 1.6 is a polynomial time recognition algorithm for determining whether G is paradox-free or not, given that G is undirected. For each $i \in [k]$, implementing Algorithm 1 with s_i in place of s and t_i in place of t , we reduce G to G_i , and obtain an answer to whether G_i is s_i - t_i series-parallel or not. If some G_i is not s_i - t_i series-parallel, then G is not paradox-free by Theorem 1.6. If G_i is s_i - t_i series-parallel for all $i \in [k]$, then the s_i - t_i block graph (see Lemma 2.8) and blocks of G_i can be identified in polynomial time for all $i \in [k]$, which implies a polynomial time examination to see if G_1, \dots, G_k satisfy the coincident condition. If the condition is satisfied, then G is paradox-free, otherwise G is paradox-ridden.

In case of G being directed, we have no efficient algorithm for determining whether an arc (u, v) is contained by an s_i - t_i path, or equivalently for determining whether the digraph has vertex-disjoint s_i - u path and v - t_i path. The latter problem is the well-known NP-complete problem of 2-vertex-disjoint directed paths (2DDP) [11], while its special case, the 2DDP problem in planar digraphs (planar 2DDP), is polynomial time solvable [7, 24]. Suppose that G is a planar digraph. Using the algorithm for planar 2DDP [7] and that for recognizing two-terminal

series-parallel digraphs [27], in polynomial time, we may identify G_i ($i \in [k]$), check whether the series-parallel and coincident conditions are satisfied by G_1, \dots, G_k , and then give the answer to whether G is paradox-free or paradox-ridden.

Corollary 5.2 *For any integer $k \geq 1$, paradox-free k -commodity undirected networks and paradox-free k -commodity planar directed networks can be recognized in polynomial time.*

5.3 Underlying Graphs v.s. Orientations

Theorem 1.6 also provides two corollaries on the relations between the paradox-freeness of a network and that of its underlying graph or orientations. Notice that if G is a directed network with origin-destination pairs $(s_i, t_i)_{i=1}^k$, then its underlying graph $\mathfrak{u}(G)$ is also a k -commodity network with the same origin-destination pairs.

Corollary 5.3 *Let G be a k -commodity directed network. If $\mathfrak{u}(G)$ is paradox-free, then so is G .*

Proof Let D be the maximum irredundant subnetwork of G with origin-destination pairs $(s_i, t_i)_{i=1}^k$. It suffices to prove that D is paradox-free. As $\mathfrak{u}(D) \subseteq \mathfrak{u}(G)$, it follows from Lemma 2.4 that $\mathfrak{u}(D)$ is paradox-free. Theorem 1.6(i) says that each $\mathfrak{u}(D)_i$ ($i \in [k]$) is s_i - t_i series-parallel. Note that $\mathfrak{u}(D_i) \subseteq \mathfrak{u}(D)_i$. It is straightforward from Definition 1.3(i) that $\mathfrak{u}(D_i)$ is s_i - t_i series-parallel for $i \in [k]$. So by Definition 1.3(ii) and Lemma 2.8, we have

- (i) for each $i \in [k]$, D_i is s_i - t_i series-parallel, and its s_i - t_i block graph is a path $v_1^i B_1^i v_2^i \cdots v_{p_i}^i B_{p_i}^i v_{p_i+1}^i$ with $v_1^i = s_i$ and $v_{p_i+1}^i = t_i$.

Next, we show that D satisfies the coincidence condition. Take arbitrary distinct $i, j \in [k]$. Suppose that B_*^i is a block of D_i and B_*^j is a block of D_j such that B_*^i and B_*^j have an arc $e = (u, v)$ in common. We need to prove $B_*^i = B_*^j$. Notice that for each $h \in \{i, j\}$, uv is an edge, written as $\mathfrak{u}(e)$, contained in $\mathfrak{u}(B_*^h) \subseteq \mathfrak{u}(D_h) \subseteq \mathfrak{u}(D)_h$. It follows from Theorem 1.6(ii) that $\mathfrak{u}(e)$ is contained in a common block $\mathfrak{u}(B)$ of $\mathfrak{u}(D)_i$ and $\mathfrak{u}(D)_j$, where B is a 2-connected subdigraph of D . Note that $\mathfrak{u}(B)$ is a two-terminal series-parallel graph that has an identical set of terminals in $\mathfrak{u}(D)_i$ and $\mathfrak{u}(D)_j$.

For each $h \in \{i, j\}$, note that $u, v \in B_*^h \cap B$. We may take m_h (resp. M_h) to be the minimum (resp. maximum) index $g \in [p_h]$ such that $B_g^h \setminus v_{g+1}^h$ (resp. $B_g^h \setminus v_g^h$) has some common vertex, say α_h (resp. β_h), with B . By (i), the choices of m_h and M_h particularly imply $D_h \cap B \subseteq \cup_{g=m_h}^{M_h} B_g^h$.

For each $h \in \{i, j\}$, considering an α_h - β_h path P in $\cup_{g=m_h}^{M_h} B_g^h$, we see that $v_g^h \in P$ for all $g = m_h + 1, \dots, M_h$. Since $P \cup B$ is 2-connected, none of $v_g^h \in P, g = m_h + 1, \dots, M_h$ can be a cut-vertex of $(\cup_{g=m_h}^{M_h} B_g^h) \cup B$. It follows from the 2-connectivity of $B_{m_h}^h, B_{m_h+1}^h, \dots, B_{M_h}^h$ and B that $(\cup_{g=m_h}^{M_h} B_g^h) \cup B$ is 2-connected. Notice that $(\cup_{g=m_h}^{M_h} \mathfrak{u}(B_g^h)) \cup \mathfrak{u}(B) \subseteq \mathfrak{u}(D)_h$ and $\mathfrak{u}(B)$ is the maximum 2-connected subgraph of

$\mathfrak{u}(D)_h$ containing $\mathfrak{u}(B)$. We deduce that $\cup_{g=m_h}^{M_h} B_g^h \subseteq B$, saying $\cup_{g=m_h}^{M_h} B_g^h \subseteq D_h \cap B$. Thus we have

- (ii) for each $h \in \{i, j\}$, $D_h \cap B = \cup_{g=m_h}^{M_h} B_g^h$ consists of a string of consecutive blocks of D_h , and in particular B_*^h is a block of $D_h \cap B$.

For each $h \in \{i, j\}$, denote $a_h = v_{m_h}^h$ and $b_h = v_{M_h}^h$. Let P_h be an s_h - t_h path in D_h with $e \in P_h$. By (i) and Lemma 2.8(ii), we see that P_h goes through $v_1^h (= s_h), v_2^h, \dots, v_{m_h}^h (= a_h), \dots, v_{M_h}^h (= b_h), \dots, v_{p_{h+1}}^h (= t_h)$ in this order. By (ii), we deduce that $P_h \cap B = P_h[a_h, b_h]$ is an a_h - b_h path. On the other hand, considering $\mathfrak{u}(P_h)$ as an s_h - t_h path in $\mathfrak{u}(D)_h$, we derive from Lemma 2.8 that $\mathfrak{u}(P_h) \cap \mathfrak{u}(B)$ is a path from the origin of $\mathfrak{u}(B)$ to the destination of $\mathfrak{u}(B)$ in $\mathfrak{u}(D)_h$. So it must be the case that

- (iii) for each $h \in \{i, j\}$, the origin and destination of $\mathfrak{u}(B)$ in $\mathfrak{u}(D)_h$ are a_h and b_h , respectively.

As $\mathfrak{u}(B)$ is a coincident block of $\mathfrak{u}(D)_i$ and $\mathfrak{u}(D)_j$, it follows from (iii) that $\{a_i, b_i\} = \{a_j, b_j\}$ is the identical terminal set of $\mathfrak{u}(B)$ in $\mathfrak{u}(D)_i$ and $\mathfrak{u}(D)_j$. Observe that $e \in P_h \cap B = P_h[a_h, b_h]$ for each $h \in \{i, j\}$. Considering $\mathfrak{u}(B)$ as an a_i - b_i series-parallel graph in $\mathfrak{u}(D)_i$, we see that $\mathfrak{u}(P_i[a_i, b_i])$ is an a_i - b_i path in B which goes through $\mathfrak{u}(e)$ in the direction from u to v .

If $a_j \neq a_i$, then $b_j = a_i, a_j = b_i$, and we find that $\mathfrak{u}(P_j[a_j, b_j]) = \mathfrak{u}(P_j[b_i, a_i])$ is an a_i - b_i path in B which goes through $\mathfrak{u}(e)$ in the opposite direction from v to u . This shows a contradiction to the fact that $\mathfrak{u}(B)$ is s_i - t_i series-parallel (recall Definition 1.3). Hence we have $a_i = a_j$ and $b_i = b_j$.

Combining $a_i = a_j$ and $b_i = b_j$ with (i) and (ii), we deduce that $D_i \cap B = D_j \cap B$, and it consists of all a_i - b_i paths (i.e., a_j - b_j paths) in B . Moreover, both B_*^i and B_*^j are blocks of $D_i \cap B = D_j \cap B$. It follows from $e \in B_*^i \cap B_*^j$ that $B_*^i = B_*^j$, which says that D satisfies the coincident condition in Theorem 1.6. Recalling (i), we deduce that D is paradox-free, as desired. □

The reverse of Corollary 5.3 is not always true. The single-commodity directed network D in Fig. 10a is paradox-free (because its maximum irredundant subnetwork

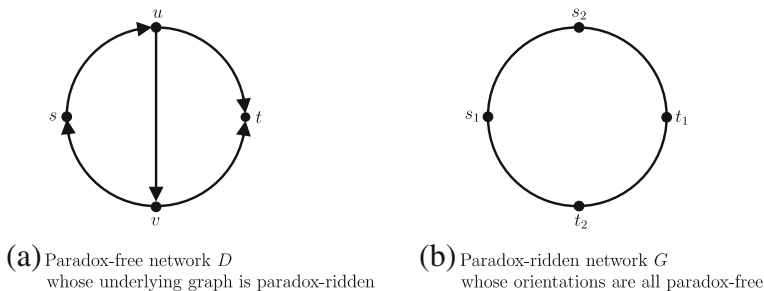


Fig. 10 Paradox-freeness concerning underlying graphs and orientations

$D \setminus (v, s)$ is s - t series-parallel), but $\mathfrak{u}(D)$ is paradox-ridden (because it is not s - t series-parallel).

On the other hand, observe that $\mathfrak{u}(D)$ can be oriented to become the digraph G_a in Fig. 1, which is paradox-ridden. Actually, this is an example for the general property of single-commodity networks stated in the following corollary. Recall Definition 3.2 of a directed s - t paradox. An undirected s - t paradox is defined in the same way via ignoring the arc directions.

Corollary 5.4 *Let G be a single-commodity undirected network with origin-destination pair (s, t) . If G is paradox-ridden, then there is an orientation \tilde{G} of G such that \tilde{G} is a paradox-ridden single-commodity directed network with origin-destination pair (s, t) .*

Proof It suffices to consider G being an irredundant network. As G is paradox-ridden, by Theorem 1.4, it is not s - t series-parallel, and thus by Proposition 1 of [20], G contains an s - t paradox (P_1, P_2, P_3) , where P_1, P_2, P_3 satisfy (i)–(iii) of Definition 3.2 in the context of graphs. Orient P_1, P_2 and P_3 into directed s - t path \tilde{P}_1 , a - v path \tilde{P}_2 and u - b path \tilde{P}_3 , respectively. Then digraph $D = \tilde{P}_1 \cup \tilde{P}_2 \cup \tilde{P}_3$ is an s - t paradox. Clearly D is an irredundant single-commodity network with origin-destination pairs (s, t) . It follows from Lemma 3.3 that D is paradox-ridden. Let \tilde{G} be a superdigraph of D . It follows from Lemma 2.4 that \tilde{G} is paradox-ridden. \square

Corollary 5.4 does not admit multicommodity extensions, as shown by the 2-commodity undirected network G in Fig. 10b. The irredundant network G does not satisfy the coincident condition, and thus by Theorem 1.6, is paradox-ridden. Consider any orientation \tilde{G} of G . The maximum irredundant subnetwork of \tilde{G} consists of an s_1 - t_1 path and an s_2 - t_2 path, and by Theorem 1.6 is paradox-free. It follows that \tilde{G} is paradox-free.

6 Conclusion

The main result of this paper is a graphical characterization for all irredundant networks to be paradox-free; the series-parallel and coincident conditions are shown to be sufficient and necessary for the paradox-freeness. Our work suggests several directions of future research.

As far as nonnegative traffic demands $\mathbf{r} = (r_i)_{i=1}^k$, nonnegative, continuous, and nondecreasing latencies ℓ are concerned, our definition of paradox-free k -commodity networks does generalize the one for single-commodity networks [17, 20, 23]. For routing instance (G, \mathbf{r}, ℓ) , we allow some traffic demands $r_i, i \in [k]$ to be zero. Clearly there is no need to consider any commodities with zero traffic demands. This is the reason why we consider the set $\mathcal{K}(\mathbf{r})$ of commodities with positive demands and subnetwork H w.r.t. the network-demand pair $(G, \mathcal{K}(\mathbf{r}))$ instead of a subnetwork defined only w.r.t. k -commodity network G itself – $\mathcal{K}(\mathbf{r})$ may be a proper subset of $[k]$, and H is a $|\mathcal{K}(\mathbf{r})|$ -commodity network in G . By manipulating \mathbf{r} , routing for (G, \mathbf{r}, ℓ) can behave as an h -commodity routing inside G for any $h \in [k]$. These

routings with $h < k$ (referred to as *degenerate routings*) are crucial for obtaining the result that any union of G_i 's ($i \in [k]$) is paradox-free provided G is paradox-free (see Lemma 2.4).

The assumption that $\mathbf{r} \in \mathbb{R}_{\geq 0}^k$ instead of $\mathbf{r} \in \mathbb{R}_{> 0}^k$ is seemingly nature because the ratio between the largest traffic demand and the smallest one can be arbitrarily large, and the arbitrarily small traffic has arbitrarily small affect on latencies. However, one might restrict \mathbf{r} to be positive; this requires each k -commodity routing for (G, \mathbf{r}, ℓ) to be nondegenerate, and each subnetwork to be a k -commodity network. With this restriction of positive demands, we would lose Lemma 2.4; it is possible that the union of some G_i 's, $i \in [k]$, is paradox-ridden even if $G = \cup_{i=1}^k G_i$ is paradox-free (see the example in Fig. 2 where $G_1 \cup G_2 \cup G_3$ in (a) is paradox-free, and $G_1 \cup G_2$ in (b) is paradox-ridden). It would be an interesting and challenging task to characterize paradox-free networks for positive traffic demands.

Different from Definition 1.2, the occurrence of Braess's paradox in a routing instance (G, \mathbf{r}, ℓ) could have another generalization from single-commodity to multicommodities as follows: We say that *strong Braess paradox occurs in* (G, \mathbf{r}, ℓ) if there exists subnetwork H of $(G, \kappa(\mathbf{r}))$ such that $\ell_i(H, \mathbf{r}) < \ell_i(G, \mathbf{r})$ for all $i \in \kappa(\mathbf{r})$. Correspondingly, we may define *strong paradox-free* networks. The class of strong paradox-free networks turns out to be much wider than that of paradox-free ones. Even under the undirected network model, the characterization of strong paradox-freeness might involve much more complicated interaction between network topologies and the positions of origin-destination pairs in the network.

A corollary of our main result is polynomial time recognition of k -commodity undirected networks and planar directed networks that are paradox-free (see Corollary 5.2). Unfortunately, this does not extend to general directed networks even restricted to the single-commodity case. The difficulty lies on the fact that we have no efficient algorithm for finding the maximum irredundant subnetwork for each single commodity, which is equivalent to solving the NP-complete problem of 2-vertex-disjoint paths. In view of this NP-completeness, the recognition of paradox-free directed networks is very likely to be intractable. It would be nice to identify the complexity status of the recognition problem.

Acknowledgments The authors are indebted to anonymous referees for their invaluable comments and suggestions which have greatly improved the presentation of this paper.

References

1. Azar, Y., Epstein, A.: The hardness of network design for unsplittable flow with selfish users. In: Erlebach, T., Persinao, G. (eds.) *Approximation and Online Algorithms*, Lecture Notes in Computer Science, vol. 3879, pp. 41–54. Springer Berlin Heidelberg (2006)
2. Beckmann, M., McGuire, B., Winsten, C.B.: *Studies in the Economics of Transportation*. Technical report (1956)
3. Bell, M.G.H., Iida, Y.: *Transportation Network Analysis* (1997)
4. Braess, D.: Über ein paradoxon aus der verkehrsplanung. *Unternehmensforschung* **12**(1), 258–268 (1968)
5. Cenciarelli, P., Gorla, D., Salvo, I.: Graph theoretic investigations on inefficiencies in network models. *Discret. Math.* (2016). arXiv:1603.01983

6. Cohen, J.E., Horowitz, P.: Paradoxical behaviour of mechanical and electrical networks. *Nature* **352**, 699–701 (1991)
7. Cygan, M., Marx, D., Pilipczuk, M.: The planar directed k -vertex-disjoint paths problem is fixed-parameter tractable. In: IEEE 54th Annual Symposium on Foundations of Computer Science (FOCS), 2013, pp. 197–206. IEEE (2013)
8. Czumaj, A.: Selfish routing on the internet. In: Leung, J.Y.-T. (ed.) *Handbook of Scheduling: Algorithms, Models, and Performance Analysis*, chap. 42. CRC Press (2004)
9. Diestel, R.: *Graph Theory*. Electronic Library of Mathematics. Springer (2000)
10. Epstein, A., Feldman, M., Mansour, Y.: Efficient graph topologies in network routing games. *Games Econ. Behav.* **66**(1), 115–125 (2009)
11. Fortune, S., Hopcroft, J., Wyllie, J.: The directed subgraph homeomorphism problem. *Theor. Comput. Sci.* **10**(2), 111–121 (1980)
12. Fotakis, D., Kaporis, A.C., Lianas, T., Spirakis, P.G.: On the hardness of network design for bottleneck routing games. *Theor. Comput. Sci.* **521**, 107–122 (2014)
13. Fujishige, S., Goemans, M.X., Harks, T., Peis, B., Zenklusen, R.: Matroids are immune to Braess paradox. *Computer Science and Game Theory* (2015). arXiv:1504.07545v1
14. Holzman, R., Monderer, D.: Strong equilibrium in network congestion games: Increasing versus decreasing costs. *Int. J. Game Theory* 1–20 (2014)
15. Holzman, R., Nissan Law yon (Lev-tov): Network structure and strong equilibrium in route selection games. *Math. Soc. Sci.* **46**(2), 193–205 (2003)
16. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. *Comput. Sci. Rev.* **3**, 65–69 (2009)
17. Lin, H., Roughgarden, T., Tardos, É., Walkover, A.: Stronger bounds on Braess’s paradox and the maximum latency of selfish routing. *SIAM J. Discret. Math.* **25**(4), 1667–1686 (2011)
18. Milchtaich, I.: Topological conditions for uniqueness of equilibrium in networks. *Math. Oper. Res.* **30**(1), 225–244 (2005)
19. Milchtaich, I.: The equilibrium existence problem in finite network congestion games. In: Spirakis, P., Mavronicolas, M., Kontogiannis, S. (eds.) *Internet and Network Economics*, Lecture Notes in Computer Science, vol. 4286, pp. 87–98. Springer Berlin Heidelberg (2006)
20. Milchtaich, I.: Network topology and the efficiency of equilibrium. *Games Econ. Behav.* **57**(2), 321–346 (2006)
21. John, D.: Murchland. Braess’s paradox of traffic flow. *Transp. Res.* **4**(4), 391–394 (1970)
22. Roughgarden, T.: On the severity of Braess’s paradox: Designing networks for selfish users is hard. *J. Comput. Syst. Sci.* **72**(5), 922–953 (2006)
23. Roughgarden, T., Tardos, É.: How bad is selfish routing? *J. ACM* **49**(2), 236–259 (2002)
24. Schrijver, A.: Finding k disjoint paths in a directed planar graph. *SIAM J. Comput.* **23**(4), 780–788 (1994)
25. Steinberg, R., Zangwill, W.I.: The prevalence of Braess’ paradox. *Transp. Sci.* **17**(3), 301–318 (1983)
26. Tutte, W.T.: *Graph Theory*. Electronic Library of Mathematics. China Machine Press (2004)
27. Valdes, J., Tarjan, R.E., Lawler, E.L.: The recognition of series parallel digraphs. *SIAM J. Comput.* **11**(2), 298–313 (1982)
28. Wardrop, J.G.: Some theoretical aspects of road traffic research. In: *Proceedings of the Institute of Civil Engineers, Part II*, vol. 1, pp. 325–378 (1952)