

On the Robustness of Potential-Based Flow Networks*

Max Klimm[†], Marc E. Pfetsch[‡], Rico Raber and Martin Skutella[§]

March 20, 2019

Abstract

Potential-based flows provide a simple yet realistic mathematical model of transport in many real-world infrastructure networks such as, e.g., electricity, gas, or water networks, where the flow along each edge is controlled via the (difference of) potentials at its end nodes. A potential-based flow network is *robust* if the maximal difference of node potentials needed to satisfy a set of demands cannot increase if demands are decreased. This notion of robustness is motivated by infrastructure networks where users first make reservations for certain demands that may be larger than the actual amounts sent later on. Here node potentials correspond to physical quantities such as the pressures or the voltages and must be guaranteed to lie within a fixed range, even if the actual amounts are smaller than the previously reserved demands. Our main results are a precise characterization of such robust networks for the case of point-to-point demands via forbidden node-labeled graph minors, as well as an efficient algorithm for testing robustness.

1 Introduction

A common feature of many infrastructure networks such as water, gas, electricity, telecommunication, and road networks is that their load heavily fluctuates due to changes in the demands of the transported commodities. As a consequence, the robustness of these networks with respect to changes in the demands is a major issue for network operators that has been studied extensively in the *robust optimization* literature; see, e.g., [2]. Informally, a network is *robust* if it is feasible for a prescribed range of demand scenarios rather than a single situation only. Naturally, the robustness of a flow network depends on the set of scenarios considered. One of the earliest robust network flow problems—the so-called *network synthesis problem*—asks for the robustness of a flow network for a given list of single-source and single-sink routing demands [9, 18, 19, 20, 25, 35]; see also [29] for a multi-commodity flow variant with discrete scenario set. Other works consider robustness of single-commodity networks with respect to a set of demands given by a full-dimensional polyhedron [1, 7, 8] or a lower dimensional polyhedron in so-called *Hose-demand sets* [8, 14, 21, 24].

*We acknowledge funding through the DFG CRC/TRR 154, Subproject A007. The last author is supported by the Einstein Foundation Berlin.

[†]School of Business and Economics, Humboldt-Universität zu Berlin, Spandauer Str. 1, 10178 Berlin, Germany, max.klimm@hu-berlin.de

[‡]Department of Mathematics, TU Darmstadt, Dolivostr. 15, 64293 Darmstadt, Germany, pfetsch@mathematik.tu-darmstadt.de

[§]Institute of Mathematics, TU Berlin, Straße des 17. Juni 136, 10623 Berlin, Germany, {raber,skutella}@math.tu-berlin.de

All these works above have in common that they consider a single commodity in the classical flow model of Ford and Fulkerson [16] with hard capacities on the edges. While obviously important (e.g., as a model for telecommunication networks) they fail to capture the flow in energy and gas distribution networks due to the more complex physical properties of the flow; see, e.g., the discussion in [8].

In addition, many infrastructure networks are controlled by a network operator who sells the right to ship a certain maximal amount of flow between a prescribed set of network nodes [22, 27]. Even in the simplest such model where network users acquire the transmission rights between pairs of distinct nodes, the set of actual demands below these maximal transmissions gives rise to a lower-dimensional set of demand scenarios that fundamentally differs from Hose scenarios and, to the best of our knowledge, has not been studied in the literature yet.

In this paper, we close both gaps, i.e., we study the robustness of infrastructure networks with respect to a natural model for physical flows under a natural robustness concept that requires robustness against all flows satisfying given maximal flow constraints. We adopt the classical potential-based flow model [5, 11, 23, 28, 31]. From a mathematical point of view, potential-based flows are particularly interesting, and also challenging, since they combine combinatorial network structures and non-linear optimization methods.

In this single-commodity model, edges are undirected and are endowed with a potential loss function that describes the flow on an edge as a function of the difference of the potentials at its endpoints. Choosing appropriate potential loss functions, potential flows model the underlying physics of stationary gas, electricity, and water networks. Specifically, a potential flow network is an undirected graph $G = (V, E)$ where each edge $e = \{u, v\}$ is endowed with a strictly increasing potential loss function $\psi_{u,v}: \mathbb{R} \rightarrow \mathbb{R}$ with $\psi_{u,v}(0) = 0$ as well as a resistance β_e (in gas networks, for example, $\psi_{u,v}(x_{u,v}) = x_{u,v} \cdot |x_{u,v}|$ [36]). Under mild assumptions on the potential loss functions, a potential vector $\pi \in \mathbb{R}^V$ induces a unique flow $x \in \mathbb{R}^{V \times V}$ via the equations

$$\beta_e \psi_{u,v}(x_{u,v}) = \pi_u - \pi_v \quad \text{for all edges } e = \{u, v\} \in E;$$

see Section 2 for details. In physical networks, the potentials correspond to physical quantities like the pressures or the voltages at the nodes. Thus, the maximal potential difference among all nodes is a measure for the stress on a network.

Routing demands are given by a finite set $(s_i, t_i)_{i \in I}$ of source-sink pairs each specifying a flow demand d_i . A flow x satisfies the demand vector $d = (d_i)_{i \in I}$ if, for all nodes u , it holds

$$\sum_{v \in V: \{u,v\} \in E} x_{u,v} = \sum_{i \in I: u=s_i} d_i - \sum_{i \in I: u=t_i} d_i.$$

Notice that the functional dependency of the potentials needed to satisfy a given set of demands is non-monotonic in the following sense: decreasing the demand d_i between some nodes s_i and t_i may increase the maximal potential difference; see Figure 1 and the discussion in Section 1.1. This is critical for the operators of infrastructures such as electricity, water, and gas networks as they usually issue the right to send maximal demands of flows between designated sets of nodes to network users [22, 27]. For safety reasons, they need to ensure that the maximal potential difference (corresponding to voltages or pressure differences in these networks) is within given bounds. Due to the non-monotonicity of the potentials, the non-fulfillment of the demand by some network user may violate the potential bounds and put the network at risk.



Figure 1: Non-robust potential flow networks.

1.1 Our results and techniques

Motivated by the discussion above, we call a potential flow network G with source-sink pairs $(s_i, t_i)_{i \in I}$ *robust* if the maximal potential difference is non-increasing as the demand decreases. That is, whenever there is a demand vector d satisfied by a flow x with maximal potential difference $\bar{\pi}$ and a demand vector d' with $d'_i \leq d_i$ for all $i \in I$, then there exists a flow x' satisfying d' with maximal potential difference $\bar{\pi}' \leq \bar{\pi}$.

In this paper, we give a full characterization of such robust networks. To give some intuition, consider the two networks depicted in Figure 1. The network in Figure 1a consists of a single edge and two routing requests in opposite directions. This network is not robust: Suppose there is a demand of one unit between s_1 and t_1 as well as another unit demand between $s_2 = t_1$ and $t_2 = s_1$. Since the requests cancel out, no flow has to be actually sent and, as $\psi_{s_1, t_1}(0) = 0$, no potential difference is needed. On the other hand, decreasing either of the demands requires an actual flow and, thus, a non-zero potential difference between the end nodes; see Lemma 8.

Moreover, consider the network in Figure 1b. A slightly more complicated reasoning shows that also this network is not robust. For ease of exposition assume that all three edges have the same potential loss function and resistance. By symmetry, to send one unit of flow from s_i to t_i for all $i \in \{1, 2, 3\}$, we can choose a potential vector π such that $\pi_{s_1} = \pi_{s_2}$ and $\pi_{t_1} = \pi_{t_3}$. On the other hand, if there is no demand between s_2 and t_2 , we need $\pi_{s_2} = \pi_{t_2}$ to prevent flow on the edge $\{s_2, t_2\}$. One can show that this inevitably leads to an increase of the maximal potential difference in the network; see Lemma 9.

As our main result, we show that these type-1 and type-2 networks in Figure 1 are essentially the only two networks that are non-robust for potential flow networks. For a formal statement of this result, we use the notion of node-labeled graph minors introduced by Friedman, Robertson, and Seymour [17] that extends the usual notion of a graph minor. To this end, we label each node with a subset of source labels s_i and sink labels t_i with $i \in I$. A graph is a minor of some other graph if it can be constructed from it by a sequence of edge contractions, edge deletions, and label deletions, where an edge contraction is defined such that the new node receives the union of the labels of its endpoints. With this definition, we show that a potential flow network is robust if and only if it does neither contain a type-1 nor a type-2 network as a (node-labeled) minor; see Theorem 14. As an immediate corollary of our result, we obtain that networks with a single source or a single sink are robust; see Corollary 16.

To exhibit the explanatory power of our characterization, we demonstrate its consequences for tree and cycle networks. Tree networks are particularly relevant since in the non-robust setting all minimal network designs are cycle-free. We show that a tree network is robust if and only if after contracting all edges that do not lie on an s_i - t_i -path, the remaining edges can be oriented such that all paths from a source s_i to a target t_i follow the orientation, and along every path in the tree the orientation of the edges flips at most once; see Theorem 24. We further give a characterization of robust cycle networks in terms of the ordering of the node labels along the cycle; see Theorem 26.

Motivated by the process of capacity nomination in the European gas market [22, 27],

we further study a situation where routing demands can specify b -vectors rather than just source-sink pairs. Specifically, we assume that the network nodes are partitioned into potential sources S and potential sinks T , and that every routing demand is a b -vector with the additional property that demands are non-positive for sources and non-negative for sinks. We show that in this model, a network is robust if and only if there is an articulation node that separates the sources from the sinks; see Theorem 30.

Finally, we give a polynomial time algorithm that determines whether a network topology is robust; see Theorem 32.

1.2 Related work

The first investigation of robust network flows is for the network synthesis problem defined by Chien [9]. Given an undirected network with demands between pairs of nodes, the problem asks for minimal edge capacities such that for each pair of nodes there is a feasible flow satisfying the demand. The problem can be reformulated as a robust flow problem in the sense of Ben-Tal and Nemirowski [3] with a discrete scenario set by introducing a scenario for each pair of vertices. There are several polynomial algorithms known for this problem [18, 20, 25, 35]. The problem also admits a linear programming formulation of polynomial size [19]. The integer version of the problem, where edge capacities are required to be integral, is studied in [10, 26, 33].

Buchheim et al. [7] consider a generalization of the integer problem where the scenario set is discrete but each scenario is an arbitrary b -flow. They show that the problem is NP-hard and propose a branch-and-cut-algorithm; see also [1, 8] for further heuristics and exact algorithms. Without the integrality constraint, the problem is solvable by linear programming techniques [32]. When the set of scenarios is a polytope, even deciding the feasibility of a point is co-NP-complete. For a multi-commodity variant of this problem, see [4, 29].

Given upper bounds on the incoming and outgoing demands for each node, the Hose polytope contains all demand matrices obeying these bounds [13, 15]. When the routing has to be fixed before the scenario is released and the scenario set is a Hose polytope, the network design problem is known as virtual private network design. The optimal solution for such a problem is always a tree [21, 24].

In recent, but unpublished work, Szabó [34] analyzed how the maximum potential difference in a potential flow network changes when inserting an additional edge. In particular it is shown that, while maintaining the demands at every node, inserting an edge may increase the stress, a phenomenon that resembles Braess's paradox [6].

2 Preliminaries

Let $G = (V, E)$ be an undirected graph. We assume that G is simple and connected. For a finite index set I , let $D = (D_i)_{i \in I}$ with $D_i = (s_i, t_i) \in V \times V$ be a set of source-sink pairs. We call the tuple (G, D) a *network topology*. A *flow* in G is a vector $x \in \mathbb{R}^{V \times V}$ with $x_{u,v} = -x_{v,u}$ for all $u, v \in V$ and $x_{u,v} = 0$ for all $u, v \in V$ with $\{u, v\} \notin E$. A positive value $x_{u,v}$ indicates a movement of flow particles along edge $\{u, v\}$ from u to v , while a negative value models flow along this edge in the opposite direction from v to u . Let

$$\mathcal{F} := \{x \in \mathbb{R}^{V \times V} : x_{u,v} = -x_{v,u} \text{ for all } u, v \in V, \\ x_{u,v} = 0 \text{ for all } u, v \in V \text{ with } \{u, v\} \notin E\}$$

denote the set of all flows in G . The *balance vector* $\text{bal}(x) \in \mathbb{R}^V$ of a flow $x \in \mathcal{F}$ is defined by $\text{bal}(x)_u := \sum_{v \in V} x_{u,v}$ for all $u \in V$. Similarly, for a vector of demands

$d \in \mathbb{R}_{\geq 0}^I$, the balance vector $\text{bal}(d) \in \mathbb{R}^V$ is defined as

$$\text{bal}(d)_u := \sum_{i \in I: u=s_i} d_i - \sum_{i \in I: u=t_i} d_i \quad \text{for all } u \in V.$$

We say that a flow x *satisfies* a vector of demands d if $\text{bal}(x) = \text{bal}(d)$.

The flows considered in this paper are based on potential vectors $\pi \in \mathbb{R}^V$. For each edge $e = \{u, v\} \in E$, we are given *potential loss functions* $\psi_{u,v}$ and $\psi_{v,u}: \mathbb{R} \rightarrow \mathbb{R}$ and a *resistance* $\beta_e \in \mathbb{R}_{>0}$. The potential loss functions $\psi_{u,v}$ and $\psi_{v,u}$ model opposite orientations of the same physical principles. Thus, we have $\psi_{u,v}(z) = -\psi_{v,u}(-z)$ for all $\{u, v\} \in E$ and $z \in \mathbb{R}$. Intuitively, the potential loss functions and resistances describe the physics of the underlying network. Recall the references in the introduction for applications of this model. Throughout this paper, we impose the following assumptions on the potential loss functions.

Assumption 1. *For each $\{u, v\} \in E$, the potential loss function $\psi_{u,v}: \mathbb{R} \rightarrow \mathbb{R}$ satisfies the following properties:*

1. $\psi_{u,v}$ is continuous,
2. $\psi_{u,v}$ is strictly increasing,
3. $\psi_{u,v}(0) = 0$.

The first two assumptions are natural in the physical context and also appear in [5]. Moreover, the third assumption is equally natural and will turn out to be essential for the characterization of robust networks.

We denote the family of all functions satisfying Assumption 1 by Ψ . We say that a flow x is induced by π if, for each edge $e = \{u, v\}$, the difference of the node potentials of the end nodes equals the potential loss induced by the flow along e , i.e.,

$$\beta_e \psi_{u,v}(x_{u,v}) = \pi_u - \pi_v. \quad (1)$$

Since the right-hand side of (1) is a difference of node potentials, uniformly shifting the entries of a potential vector $\pi \in \mathbb{R}^V$ has no effect in terms of (1). Since the potential loss functions $\psi_{u,v}$ are one-to-one, for a given potential vector $\pi \in \mathbb{R}^V$ with $(\pi_u - \pi_v)/\beta_e \in \psi_{u,v}(\mathbb{R})$ for all $e = \{u, v\} \in E$, there is a unique flow $x(\pi) \in \mathcal{F}$ satisfying (1). For such a flow $x(\pi)$, the corresponding balance vector $\text{bal}(x(\pi))$ can be computed as

$$\text{bal}(x(\pi))_u = \sum_{v \in V} x_{u,v}(\pi) = \sum_{v \in V: \{u,v\} \in E} \psi_{u,v}^{-1} \left(\frac{\pi_u - \pi_v}{\beta_{\{u,v\}}} \right) \quad (2)$$

for all $u \in V$.

Let $B := \{b \in \mathbb{R}^V : \sum_{v \in V} b_v = 0\}$ be the set of node balances that sum to 0, and let

$$\Pi_{v_0} := \left\{ \pi \in \mathbb{R}^V : (\pi_u - \pi_v)/\beta_{\{u,v\}} \in \psi_{u,v}(\mathbb{R}) \text{ for all } \{u, v\} \in E, \pi_{v_0} = 0 \right\}$$

be the set of potential vectors for which (without loss of generality) the potential of some node $v_0 \in V$ is fixed to 0. Under the above conditions on G and the potential loss functions $\psi_{u,v}$, for all $v_0 \in V$, the function $f: \Pi_{v_0} \rightarrow B$ defined as $f(\pi) := (f_u(\pi))_{u \in V}$ with

$$f_u(\pi) := \sum_{v \in V: \{u,v\} \in E} \psi_{u,v}^{-1} \left(\frac{\pi_u - \pi_v}{\beta_{\{u,v\}}} \right) \quad (3)$$

is bijective and continuous. In particular, the inverse function $f^{-1}: B \rightarrow \Pi_{v_0}$ exists and is also continuous; see, e.g., Birkhoff and Diaz [5]. This implies that there is a one-to-one correspondence between node balances $b \in B$ and potentials $\pi \in \Pi_{v_0}$.

An important measure is the *stress* on a network G which is defined as the maximal potential difference of the nodes; see, e.g., [31, 34], where the influence of modifications of the underlying network on the stress is analyzed. For a potential vector $\pi \in \Pi_{v_0}$, we write

$$\text{str}_G(\pi) := \max_{v \in V} \pi_v - \min_{v \in V} \pi_v.$$

Using a slight overload of notation, we define for a balance vector $b \in B$ the stress $\text{str}_G(b)$ of the corresponding potential vector $\pi = f^{-1}(b)$ as

$$\text{str}_G(b) := \max_{v \in V} (f^{-1}(b))_v - \min_{v \in V} (f^{-1}(b))_v.$$

Further overloading the notation, we write for a demand vector $d \in \mathbb{R}_{\geq 0}^I$

$$\text{str}_G(d) := \text{str}_G(\text{bal}(d)).$$

It is straightforward to see that the stress of a balance vector $b \in B$ is invariant under the choice of v_0 . If the network G is clear from the context, we omit the subscript and simply write $\text{str}(\pi)$, $\text{str}(b)$, and $\text{str}(d)$.

3 Robust networks

In this section we introduce the concept of *robustness* of a network topology and give a full characterization of robust network topologies. We call a network topology *robust* if a component-wise decrease of a demand vector d never leads to an increase of the stress.

Definition 1. A network topology (G, D) together with potential loss functions $\psi_{u,v} \in \Psi$, for $\{u, v\} \in E$, is called *robust* if, for all $\beta \in \mathbb{R}_{> 0}^E$, the function $\text{str}: \mathbb{R}_{\geq 0}^I \rightarrow \mathbb{R}_{\geq 0}$ is non-decreasing, i.e., for all $d', d \in \mathbb{R}_{\geq 0}^I$ with $d' \leq d$ (component-wise), we have $\text{str}(d') \leq \text{str}(d)$.

Remark 2. One may also want to consider a stronger form of robustness, where the monotonicity of the stress even holds for *all* potential loss functions $\psi_{u,v}$, $\{u, v\} \in E$. We call a network topology *strongly robust* if, for all $\beta \in \mathbb{R}_{> 0}^E$ and for all $\psi_{u,v} \in \Psi$, $\{u, v\} \in E$, the function $\text{str}: \mathbb{R}_{\geq 0}^I \rightarrow \mathbb{R}_{\geq 0}$ is non-decreasing. As a byproduct of our analysis below, we prove that a network topology is robust if and only if it is strongly robust.

For our characterization of robust network topologies, we need the concept of a *minor* of a network topology, which is a node-labeled graph minor in which the labels form a quasi-order as introduced by Friedman, Robertson, and Seymour [17]. To this end, consider a node-labeled graph G_L with label set $L := \{\mathfrak{s}_i : i \in I\} \cup \{\mathfrak{t}_i : i \in I\}$, and define $\ell(v) \subseteq L$ to be the subset of labels attached to node $v \in V$. The graph is *well-labeled* if, for all $i \in I$,

$$|\{v \in V : \mathfrak{s}_i \in \ell(v)\}| = |\{v \in V : \mathfrak{t}_i \in \ell(v)\}| \in \{0, 1\},$$

i.e., labels \mathfrak{s}_i and \mathfrak{t}_i are used pairwise or not at all.

There is a bijection between network topologies and well-labeled graphs: For a given network topology (G, D) , each node v obtains the label set $\ell(v) = \{\mathfrak{s}_i : i \in I \text{ with } s_i =$




Figure 2: Original network topology (G, D) and one of its minors (\bar{G}, \bar{D}) . A possible sequence of contractions and deletions is as follows: contraction of edge e_7 , deletion of edge e_4 , deletion of edge e_2 , contraction of edge e_1 and deletion of labels s_1, t_1 .

$v\} \cup \{t_i : i \in I \text{ with } t_i = v\}$. Conversely, each well-labeled graph defines a network topology using those node pairs (s_i, t_i) with $s_i \in \ell(s_i)$ and $t_i \in \ell(t_i)$.

For a well-labeled graph, the contraction of an edge $e = \{u, w\} \in E$ is the operation that deletes e , merges u and w into a single node, and gives this node the label set $\ell(u) \cup \ell(w)$. We can delete a label pair s_i, t_i by deleting the labels s_i and t_i from both L and the label sets they are contained in, and deleting i from I . Deletion of edges is defined analogously to the unlabeled case. A labeled graph \bar{G}_L is a *minor* of a labeled graph G_L , if the former can be constructed from the latter by a finite sequence of edge contractions, edge deletions, and label deletions.

Definition 3. Let (G, D) be a network topology and G_L the corresponding well-labeled graph. Then a network topology (\bar{G}, \bar{D}) is a minor of (G, D) if its corresponding well-labeled graph \bar{G}_L is a minor of G_L and if \bar{G} is connected and simple.

Figure 2 depicts an example of a minor of a network topology. The facilitating requirement in Definition 3 that a minor must be connected and simple is, for our purposes, without loss of generality as the two minors occurring in our main result are indeed connected and simple; see Figure 1. Due to their one-to-one correspondence, throughout this paper we use the notions of *network topologies* and *well-labeled graphs* interchangeably.

If (\bar{G}, \bar{D}) is a minor of (G, D) , then, in particular, \bar{G} is an (ordinary unlabeled) minor of G ; see, e.g., Diestel [12]. In particular, any node in a minor corresponds to a connected subgraph in the original graph, and any two such connected subgraphs are disjoint, a fact that turns out to be useful later.

Lemma 4 ([12, Section 1.7]). *Let (G, D) be a network topology and (\bar{G}, \bar{D}) a minor, where $G = (V, E)$ and $\bar{G} = (\bar{V}, \bar{E})$. For $\bar{v} \in \bar{V}$, let $V(\bar{v}) \subseteq V$ be the subset of nodes in V contracted into \bar{v} when creating the minor. Then, for every $\bar{v} \in \bar{V}$, the induced subgraph $G[V(\bar{v})]$ is connected, and for any two $\bar{v}_1, \bar{v}_2 \in \bar{V}$ with $\bar{v}_1 \neq \bar{v}_2$, we have $V(\bar{v}_1) \cap V(\bar{v}_2) = \emptyset$.*

The sequence of operations to obtain a minor (\bar{G}, \bar{D}) of (G, D) can always be chosen such that all intermediate graphs are connected and simple. This follows from the observation that, whenever an edge contraction results in a loop, one can instead delete the loop edge before the contraction. Moreover, no intermediate graph can be disconnected as otherwise also \bar{G} is disconnected. As a consequence, when constructing a minor, we only need to consider the following *basic operations*.

Definition 5. For a network topology (G, D) the following are *basic operations*:

- (a) deletion of an edge which lies on a cycle;
- (b) contraction of an edge which does not lie on a triangle (cycle of length three);
- (c) deletion of a pair of labels $\mathfrak{s}_i, \mathfrak{t}_i$.

We show that every minor of a robust network topology (G, D) is robust. In order to prove this result, we need the following lemma. It states that, for each of the three operations above, we can adapt the vector of edge resistances β such that the stress on the original network and the stress on the minor network are arbitrarily close to each other. The intuition behind the proof is that deletion or contraction of an edge can be approximated by giving the edge a very high or very low resistance, respectively. For given node potentials, this approximation distorts the node balances only by a very small amount. Then, by continuity, one only needs to change the potentials (and thus, in particular, the stress on the network) by a very small amount to restore the original node balances.

Lemma 6. *Let (G, D) be a network topology and (\bar{G}, \bar{D}) with $\bar{G} = (\bar{V}, \bar{E})$ and $\bar{D} = (\bar{D}_i)_{i \in \bar{I}}$ a minor obtained by one basic operation. Then, for every $\bar{d} \in \mathbb{R}^{\bar{I}}$, $\bar{\beta} \in \mathbb{R}_{>0}^{\bar{E}}$, and $\varepsilon > 0$, there exists $\beta \in \mathbb{R}_{>0}^E$ with $\beta_e = \bar{\beta}_e$ for all $e \in \bar{E} \subseteq E$ such that*

$$|\text{str}_G(d) - \text{str}_{\bar{G}}(\bar{d})| < \varepsilon,$$

where $d \in \mathbb{R}_{\geq 0}^V$ is defined as $d_i := \bar{d}_i$ for all $i \in \bar{I}$ and $d_i := 0$ for all $i \in I \setminus \bar{I}$.

Proof. The proof is trivial for the case that the basic operation deletes a pair of labels, so we only discuss the remaining two cases. Let $\bar{d} \in \mathbb{R}^{\bar{I}}$, $\bar{\beta} \in \mathbb{R}_{>0}^{\bar{E}}$, and $\varepsilon > 0$. Let $e^* \in E \setminus \bar{E}$ be the edge which is deleted or contracted by the basic operation. Let $\beta \in \mathbb{R}^E$ with $\beta_e = \bar{\beta}_e$ for all $e \in \bar{E}$ and β_{e^*} arbitrary. The value of β_{e^*} will be determined later, depending on the basic operation performed to produce the minor. Let $v_0 \in V \cap \bar{V}$, and let $f: \Pi_{v_0} \rightarrow B$ and $\bar{f}: \bar{\Pi}_{v_0} \rightarrow \bar{B}$ be the functions mapping potentials to balances as defined in (3) for the original graph G and its minor \bar{G} , respectively. Let $\pi := f^{-1}(\text{bal}(d))$ and $\bar{\pi} := \bar{f}^{-1}(\text{bal}(\bar{d}))$ be the potentials corresponding to d and \bar{d} , respectively. We show that one can choose β_{e^*} such that

$$|\text{str}_G(d) - \text{str}_{\bar{G}}(\bar{d})| = |\text{str}_G(\pi) - \text{str}_{\bar{G}}(\bar{\pi})| < \varepsilon.$$

We distinguish two cases depending on the conducted basic operation on e^* .

First case: e^* is deleted (and thus lies on a cycle). Then, $\bar{E} = E \setminus \{e^*\}$ and $\bar{V} = V$ and, by construction of d , $\text{bal}(d) = \text{bal}(\bar{d})$. For the potential vector $\bar{\pi} = \bar{f}^{-1}(\text{bal}(\bar{d}))$, inserting edge $e^* = \{u, w\}$ into graph \bar{G} only changes the balances of nodes u and w , namely by the amount of flow along edge e^* . For $b := f(\bar{\pi})$ we get

$$\|b - \text{bal}(d)\|_\infty = \|b - \text{bal}(\bar{d})\|_\infty = \|f(\bar{\pi}) - \bar{f}(\bar{\pi})\|_\infty = \left| \psi_{u,w}^{-1} \left(\frac{\bar{\pi}_u - \bar{\pi}_w}{\beta_{e^*}} \right) \right|.$$

Since $\psi_{u,w}$ is continuous with continuous inverse and $\psi_{u,w}(0) = 0$, we obtain that $\|b - \text{bal}(d)\|_\infty \rightarrow 0$ for $\beta_{e^*} \rightarrow \infty$. Since f^{-1} is continuous as well, one can choose β_{e^*} large enough such that

$$\|\pi - \bar{\pi}\|_\infty = \|f^{-1}(\text{bal}(d)) - f^{-1}(b)\|_\infty < \frac{1}{2}\varepsilon,$$

and hence $|\text{str}_G(\pi) - \text{str}_{\bar{G}}(\bar{\pi})| < \varepsilon$.

Second case: e^* is contracted (and thus does not lie on a triangle). For $e^* = \{u, w\}$, nodes u and w are contracted into a single node denoted by v^* . We have $\text{bal}(d)_v = \text{bal}(\bar{d})_v$ for all $v \in V \setminus \{u, w\}$ and $\text{bal}(\bar{d})_{v^*} = \text{bal}(d)_u + \text{bal}(d)_w$. Consider the potential vector $\pi^1 \in \mathbb{R}^V$ defined as

$$\pi_v^1 := \begin{cases} \bar{\pi}_{v^*} & \text{if } v \in \{u, w\}, \\ \bar{\pi}_v & \text{otherwise,} \end{cases}$$

and note that

$$\text{str}_G(\pi^1) = \text{str}_{\bar{G}}(\bar{\pi}). \quad (4)$$

Let $b^1 := f(\pi^1)$. Then $b_v^1 = \text{bal}(d)_v$ for all $v \in V \setminus \{u, w\}$ and $b_u^1 + b_w^1 = \text{bal}(\bar{d})_{v^*} = \text{bal}(d)_u + \text{bal}(d)_w$. It is without loss of generality to assume that $b_w^1 \leq \text{bal}(d)_w$ and, thus, $b_u^1 \geq \text{bal}(d)_u$. In order to restore the balances at u and w , we send a flow of value $\text{bal}(d)_w - b_w^1$ from w to u along edge $e^* = \{u, w\}$, by decreasing the potential at u . To this end, let $\pi^2 \in \mathbb{R}^V$ be defined as

$$\pi_v^2 := \begin{cases} \pi_u^1 - \beta_{e^*} \psi_{w,u}(\text{bal}(d)_w - b_w^1) & \text{if } v = u, \\ \pi_v^1 & \text{otherwise.} \end{cases}$$

Note that

$$\pi^2 \rightarrow \pi^1 \quad \text{for } \beta_{e^*} \rightarrow 0. \quad (5)$$

Let $b^2 := f(\pi^2)$. By construction, we have

$$b_w^2 = \text{bal}(d)_w \quad \text{and} \quad b_v^2 = \text{bal}(d)_v \quad \text{for all } v \in V \setminus (N(u) \cup \{u\}), \quad (6)$$

where $N(u) := \{v \in V : \{u, v\} \in E\}$ is the set of neighbors of node u . However, by having decreased the potential at u , the balance of all neighbors of u has increased. For $v \in N(u) \setminus \{w\}$ we have

$$\begin{aligned} b_v^2 - \text{bal}(d)_v &= b_v^2 - b_v^1 \\ &= \psi_{v,u}^{-1} \left(\frac{\pi_v^2 - \pi_u^2}{\beta_{\{v,u\}}} \right) - \psi_{v,u}^{-1} \left(\frac{\pi_v^1 - \pi_u^1}{\beta_{\{v,u\}}} \right) \stackrel{(5)}{\rightarrow} 0 \quad \text{for } \beta_{e^*} \rightarrow 0, \end{aligned} \quad (7)$$

by continuity of $\psi_{v,u}^{-1}$. Furthermore, since

$$\sum_{v \in V} (b_v^2 - \text{bal}(d)_v) = \sum_{v \in V} b_v^2 - \sum_{v \in V} \text{bal}(d)_v = 0 - 0 = 0,$$

we have

$$\begin{aligned} b_u^2 - \text{bal}(d)_u &= - \sum_{v \in V \setminus \{u\}} (b_v^2 - \text{bal}(d)_v) \\ &\stackrel{(6)}{=} - \sum_{v \in N(u) \setminus \{w\}} (b_v^2 - \text{bal}(d)_v) \stackrel{(7)}{\rightarrow} 0 \quad \text{for } \beta_{e^*} \rightarrow 0. \end{aligned} \quad (8)$$

Equations (6), (7), and (8) imply $b^2 - \text{bal}(d) \rightarrow 0$ as $\beta_{e^*} \rightarrow 0$ and, hence,

$$\pi - \pi^2 = f^{-1}(\text{bal}(d)) - f^{-1}(b^2) \rightarrow 0 \quad \text{for } \beta_{e^*} \rightarrow 0, \quad (9)$$

by continuity of f^{-1} . Altogether, by (4), (5), and (9), we get for β_{e^*} small enough

$$\begin{aligned} &|\text{str}_G(\pi) - \text{str}_{\bar{G}}(\bar{\pi})| \\ &= |\text{str}_G(\pi) - \text{str}_G(\pi^2) + \text{str}_G(\pi^2) - \text{str}_G(\pi^1) + \text{str}_G(\pi^1) - \text{str}_{\bar{G}}(\bar{\pi})| < \varepsilon, \end{aligned}$$

which completes the proof. \square

We are now in the position to show that robustness of a network topology is closed under the minor operation.

Lemma 7. *Every minor of a robust network topology is robust.*

Proof. Let (G, D) be a robust network topology. By contradiction, assume there is a non-robust minor (\bar{G}, \bar{D}) with $\bar{G} = (\bar{V}, \bar{E})$ and $\bar{D} = (\bar{D}_i)_{i \in \bar{I}}$. Since (\bar{G}, \bar{D}) is obtained by a sequence of basic operations, by considering the last minor in the sequence that is robust, we may assume that (\bar{G}, \bar{D}) is obtained from (G, D) by one basic operation.

As (\bar{G}, \bar{D}) is not robust, there are $\bar{\beta} \in \mathbb{R}_{\geq 0}^{\bar{E}}$ and $\bar{d} \leq \bar{d}' \in \mathbb{R}_{\geq 0}^{\bar{I}}$ for which

$$\text{str}_{\bar{G}}(\bar{d}) \geq \text{str}_{\bar{G}}(\bar{d}') + \varepsilon \quad (10)$$

for some $\varepsilon > 0$.

For the network topology (G, D) with $G = (V, E)$ and $D = (D_i)_{i \in I}$ consider the demand vectors $d, d' \in \mathbb{R}_{\geq 0}^I$ with $d_i = \bar{d}_i, d'_i = \bar{d}'_i$ for all $i \in \bar{I}$ and $d_i = d'_i = 0$ for all $i \in I \setminus \bar{I}$. By construction, $d \leq d'$. Lemma 6 implies the existence of $\beta \in \mathbb{R}^E$ with $\beta_e = \bar{\beta}_e$ for all $e \in \bar{E} \subseteq E$ such that $|\text{str}_G(d) - \text{str}_{\bar{G}}(\bar{d})| < \varepsilon/2$ and $|\text{str}_G(d') - \text{str}_{\bar{G}}(\bar{d}')| < \varepsilon/2$. Due to (10) it follows that $\text{str}_G(d) > \text{str}_G(d')$, contradicting the robustness of (G, D) . \square

We proceed to show that a network topology is robust if and only if it contains neither of two special minors, called *type-1* and *type-2* networks depicted in Figure 1. Before we prove the full characterization of robustness, we show that type-1 and type-2 network topologies are not robust.

Lemma 8. *The type-1 network topology is not robust.*

Proof. Consider the network in Figure 1a with a demand vector d such that $d_1 = d_2 > 0$. Then there is no flow on the only edge $e = \{s_1, t_1\}$, and since $\psi_{s_1, t_1}(0) = 0$, it follows that $\text{str}(d) = 0$. Consider the demand vector d' with $d'_1 = d_1$ and $d'_2 = 0$, such that $d' \leq d$. Since $\text{bal}(d')_{s_1} = -\text{bal}(d')_{t_1} = d_1 > 0$, $\psi_{s_1, t_1}(0) = 0$, and ψ_{s_1, t_1} is strictly increasing, a positive potential difference between the two nodes is necessary in order to enforce a flow, which implies $\text{str}(d') > 0$. \square

Lemma 9. *The type-2 network topology is not robust.*

Proof. Consider the type-2 network in Figure 3a. Let $\beta_{e_1} := 1/\psi_{v_1, v_2}(1)$, $\beta_{e_2} := 1/\psi_{v_3, v_2}(1)$, $\beta_{e_3} := 1/\psi_{v_3, v_4}(1)$, and consider the demand vector d with $d_i = 1$ for $i = 1, 2, 3$. Let $\pi := f^{-1}(\text{bal}(d))$. Then $\text{str}(\pi) = 1$; see Figure 3b. The flow x induced by π satisfies $x_{v_1, v_2} = x_{v_3, v_4} = 1$ and $x_{v_2, v_3} = -1$. Next, consider the demand vector d' with $d'_1 = d'_3 = 1$ and $d'_2 = 0$. Let $\pi' := f^{-1}(\text{bal}(d'))$. Then $\text{str}(\pi') = 2$; see Figure 3c. (Note that the flow x' induced by π' satisfies $x'_{v_1, v_2} = x'_{v_3, v_4} = 1$ and $x'_{v_2, v_3} = 0$.) Hence, the type-2 network is not robust. \square

We proceed to characterize the network topologies that have a type-1 network as a minor. To this end, we introduce the following notation. For an undirected graph $G = (V, E)$, a *path* P is a sequence of pairwise distinct nodes (v_1, \dots, v_k) such that $\{v_i, v_{i+1}\} \in E$ for all $i \in \{1, \dots, k-1\}$. We denote by $V(P) := \{v_1, \dots, v_k\}$ and $E(P) := \{\{v_i, v_{i+1}\} : i \in \{1, \dots, k-1\}\}$ the node and edge set of P , respectively. Two paths P and P' are called *node-disjoint* if $V(P) \cap V(P') = \emptyset$. For two nodes $u, v \in V$, we call a path (v_1, \dots, v_k) a *u-v-path* if $v_1 = u$ and $v_k = v$. We denote the set of all *u-v-paths* in G by $\mathcal{P}_{u,v}^G$.

Furthermore, for a network topology (G, D) and two labels $\mathbf{u}, \mathbf{v} \in L$ (of the corresponding well-labeled graph G_L), a *u-v-path* is a *u-v-path*, where u and v are the nodes labeled by \mathbf{u} and \mathbf{v} , respectively, i.e., $\mathbf{u} \in \ell(u)$ and $\mathbf{v} \in \ell(v)$. Correspondingly, we set $\mathcal{P}_{\mathbf{u}, \mathbf{v}}^G := \mathcal{P}_{u,v}^G$. If the graph G is clear from the context, we sometimes omit the superscript and simply write $\mathcal{P}_{\mathbf{u}, \mathbf{v}}$ and $\mathcal{P}_{u,v}$.




Figure 3: The type-2 network with node potentials satisfying demands d and d' .




Figure 4: Node-disjoint paths $P \in \mathcal{P}_{s_i, t_j}$ and $P' \in \mathcal{P}_{s_j, t_i}$ imply the existence of a type-1 minor.

Lemma 10. *Let (G, D) be a network topology, (\bar{G}, \bar{D}) a minor, and $u_1, v_1, u_2, v_2 \in \bar{L} \subseteq L$ labels. If there exist two node-disjoint paths $\bar{P}_1 \in \mathcal{P}_{u_1, v_1}^{\bar{G}}$ and $\bar{P}_2 \in \mathcal{P}_{u_2, v_2}^{\bar{G}}$ in the minor \bar{G} , then there also exist two node-disjoint paths $P_1 \in \mathcal{P}_{u_1, v_1}^G$ and $P_2 \in \mathcal{P}_{u_2, v_2}^G$ in G .*

Proof. Let $\bar{P}_1 \in \mathcal{P}_{u_1, v_1}^{\bar{G}}$ and $\bar{P}_2 \in \mathcal{P}_{u_2, v_2}^{\bar{G}}$ be node-disjoint paths. Denote the node sets of \bar{P}_1 and \bar{P}_2 by \bar{V}_1 and \bar{V}_2 , respectively. By Lemma 4, $V_1 := \bigcup_{\bar{v} \in \bar{V}_1} V(\bar{v}) \subset V$ and $V_2 := \bigcup_{\bar{v} \in \bar{V}_2} V(\bar{v}) \subset V$ are disjoint, and both $G[V_1]$ and $G[V_2]$ are connected. Furthermore, $u_1, v_1 \in \bigcup_{v \in V_1} \ell(v)$ and $u_2, v_2 \in \bigcup_{v \in V_2} \ell(v)$. Hence, there exist node-disjoint paths $P_1 \in \mathcal{P}_{u_1, v_1}^G$ and $P_2 \in \mathcal{P}_{u_2, v_2}^G$ in G . \square

Lemma 11. *A network topology (G, D) contains a type-1 minor if and only if there exist $i, j \in I$ with $i \neq j$ and two node-disjoint paths $P \in \mathcal{P}_{s_i, t_j}^G$ and $P' \in \mathcal{P}_{s_j, t_i}^G$.*

Proof. “ \Leftarrow ”: Suppose there are $i, j \in I$ with $i \neq j$ and two node-disjoint paths $P \in \mathcal{P}_{s_i, t_j}^G$ and $P' \in \mathcal{P}_{s_j, t_i}^G$. Since G is connected, there exists a path Q connecting P and P' with $|V(Q) \cap V(P)| = |V(Q) \cap V(P')| = 1$; see Figure 4a. Consider the minor (\bar{G}, \bar{D}) which is obtained from (G, D) as follows: all edges in $E \setminus (E(P) \cup E(P') \cup E(Q))$ are deleted and all edges in $E(P) \cup E(P')$ and all edges in $E(Q)$ except for a single one are contracted. Further, all labels except $s_i, t_i, s_j,$ and t_j are deleted. This yields a type-1 network; see Figure 4b.

“ \Rightarrow ”: Let (G, D) contain a type-1 minor (\bar{G}, \bar{D}) . The type-1 network obviously contains node-disjoint-paths $\bar{P} \in \mathcal{P}_{s_i, t_j}^{\bar{G}}$ and $\bar{P}' \in \mathcal{P}_{s_j, t_i}^{\bar{G}}$, namely paths consisting only of a single node. Therefore, by Lemma 10, also the original network contains node-disjoint paths $P \in \mathcal{P}_{s_i, t_j}^G$ and $P' \in \mathcal{P}_{s_j, t_i}^G$. \square

In order to prove the general characterization of robust networks we need some preparation. For a graph $G = (V, E)$ and $s, t \in V$ with $s \neq t$, we call a flow $x \in \mathcal{F}$ an s - t -flow if $\text{bal}(x)_s = -\text{bal}(x)_t \geq 0$ and $\text{bal}(x)_v = 0$ for all $v \in V \setminus \{s, t\}$. Furthermore, for two disjoint subsets $U, W \subset V$, let $[U, W] := \{\{u, w\} \in E : u \in U, w \in W\}$ be the cut between U and W .

Lemma 12. Let $G = (V, E)$ be a connected graph, $x \in \mathcal{F}$ an s - t -flow, and $\pi \in \mathbb{R}^V$ such that

$$\operatorname{sgn}(x_{u,v}) = \operatorname{sgn}(\pi_u - \pi_v) \quad \text{for all } \{u, v\} \in E. \quad (11)$$

Then the following holds:

- (a) $\pi_s \geq \pi_v \geq \pi_t$ for all $v \in V$.
- (b) The flow x can be decomposed into a sum of positive flows along a set of s - t -paths $\mathcal{P} \subseteq \mathcal{P}_{s,t}$, such that for every path $(v_1, \dots, v_k) \in \mathcal{P}$ we have $x_{v_i, v_{i+1}} > 0$ for all $i \in \{1, \dots, k-1\}$.
- (c) For all $c \in \mathbb{R}$, both the subgraph induced by $V_c^+ := \{v \in V : \pi_v \geq c\}$ and the subgraph induced by $V_c^- := \{v \in V : \pi_v < c\}$ are connected.

Proof. We first show (a). Let $\bar{\pi} := \max_{v \in V} \pi_v$ and assume by contradiction that $\pi_s < \bar{\pi}$. Let $\bar{V} := \{v \in V : \pi_v = \bar{\pi}\}$. Since $s \notin \bar{V}$ and G is connected, there is $\{u, v\} \in E$ with $u \in \bar{V}$ and $v \in V \setminus \bar{V}$. By the maximality of π_u and (11), we have $x_{u,w} \geq 0$ for all $w \in V$ and $x_{u,v} > 0$. Therefore,

$$\operatorname{bal}(x)_u = \sum_{w \in V} x_{u,w} \geq x_{u,v} > 0,$$

contradicting the fact that x is an s - t -flow. Similarly, one can conclude that $\pi_t = \min_{v \in V} \pi_v$.

To show (b), note that by classical flow decomposition the s - t -flow x decomposes into a sum of positive flows along a set of s - t -paths \mathcal{P} and a set of cycles \mathcal{C} , such that all edges carry positive flow in the direction of the path or cycle. Assume by contradiction that $\mathcal{C} \neq \emptyset$ and $(v_1, \dots, v_k, v_1) \in \mathcal{C}$. Then (11) implies $\pi_{v_1} > \pi_{v_2} > \dots > \pi_{v_k} > \pi_{v_1}$, a contradiction.

To prove (c), we only need to consider values c with $\min_{v \in V} \pi_v < c \leq \max_{v \in V} \pi_v$, since V_c^+ or V_c^- is empty otherwise. Let $C := [V_c^+, V_c^-] \subseteq E$ be the cut between V_c^+ and V_c^- . Consider $u, v \in V_c^+$ and $P \in \mathcal{P}_{u,v}$. We are done if all nodes of P are contained in V_c^+ . Otherwise, P contains at least two edges of the cut C . Let $e_\alpha = \{\alpha^+, \alpha^-\}$ and $e_\omega = \{\omega^+, \omega^-\}$ with $\alpha^+, \omega^+ \in V_c^+$ and $\alpha^-, \omega^- \in V_c^-$ be the first and the last edge of P contained in C . For $\sigma \in \{\alpha, \omega\}$ we have $\pi_{\sigma^+} > \pi_{\sigma^-}$ and thus $x_{\sigma^+, \sigma^-} > 0$. By (b) this implies the existence of an s - t -path $P_\sigma \in \mathcal{P}$ containing edge $\{\sigma^+, \sigma^-\}$. The path P_σ does not contain any other edge of C , since all edges in C carry positive flow from V_c^+ to V_c^- . Therefore, P_σ contains a subpath from s to σ^+ whose nodes are all contained in V_c^+ . It follows that there exists a path from u to α^+ , from α^+ to s , from s to ω^+ , and from ω^+ to v only using nodes in V_c^+ . Thus, the subgraph induced by V_c^+ is connected. Analogously, one can show that the subgraph induced by V_c^- is connected. \square

The following lemma gives a necessary condition on the potential vectors of flows in type-1-free networks.

Lemma 13. Let (G, D) be a network topology without a type-1 minor, let $x \in \mathcal{F}$ be an s_j - t_j -flow for some $j \in I$, and let $\pi \in \mathbb{R}^V$ fulfill (11). Then, $\pi_{s_i} \geq \pi_{t_i}$ for all $i \in I$.

Proof. By contradiction, suppose that $\pi_{t_i} > \pi_{s_i}$ for some $i \in I \setminus \{j\}$. By Lemma 12a it follows that $\pi_{s_j} \geq \pi_{t_i} > \pi_{s_i} \geq \pi_{t_j}$. Let $V^+ := \{v \in V : \pi_v \geq \pi_{t_i}\}$ and $V^- := V \setminus V^+$. Then, $s_j, t_i \in V^+$ and $t_j, s_i \in V^-$. By Lemma 12c, both $G[V^+]$ and $G[V^-]$ are connected. Therefore, there exist node-disjoint paths $P \in \mathcal{P}_{s_i, t_j}$ and $P' \in \mathcal{P}_{s_j, t_i}$. By Lemma 11, (G, D) contains a type-1 minor, a contradiction. \square

We can now prove the main result of this section.

Theorem 14. *A network topology (G, D) is robust if and only if it does neither contain a type-1 nor a type-2 minor.*

Proof. “ \Rightarrow ”: If (G, D) contains a type-1 or type-2 minor, then, by Lemmas 8 and 9, (G, D) has a minor that it not robust. Lemma 7 implies that (G, D) is not robust.

“ \Leftarrow ”: For the reverse direction assume that (G, D) is not robust, i.e., there exist $\beta \in \mathbb{R}_{>0}^E$ and $d, d' \in \mathbb{R}_{\geq 0}^I$ with $d \leq d'$ such that $\text{str}(d') < \text{str}(d)$. We can assume without loss of generality that there is $j \in I$ such that $d_j < d'_j$ and $d_i = d'_i$ for all $i \in I \setminus \{j\}$; otherwise, move from d' to d by successively decreasing the demands d'_i to d_i , one at a time, and identify a step which strictly increases the stress.

Let $\pi := f^{-1}(\text{bal}(d))$, $\pi' := f^{-1}(\text{bal}(d'))$, and let $x := x(\pi)$, $x' := x(\pi')$ be the corresponding flows. Then, $\Delta x := x' - x$ is an s_j - t_j -flow of value $d'_j - d_j > 0$. Due to Equation (1) and Assumption 1, the flow Δx and the vector of potential differences $\Delta\pi := \pi' - \pi$ satisfy

$$\begin{aligned} \text{sgn}(\Delta x_{u,v}) &= \text{sgn}(x'_{u,v} - x_{u,v}) \\ &= \text{sgn}(\psi_{u,v}^{-1}(\pi'_u - \pi'_v) - \psi_{u,v}^{-1}(\pi_u - \pi_v)) \\ &= \text{sgn}((\pi'_u - \pi'_v) - (\pi_u - \pi_v)) \\ &= \text{sgn}(\Delta\pi_u - \Delta\pi_v) \quad \text{for all } \{u, v\} \in E. \end{aligned} \tag{12}$$

Let V_{\max} and V_{\min} be the node sets of an arbitrary connected component of $G[\text{argmax}_{v \in V} \pi_v]$ and $G[\text{argmin}_{v \in V} \pi_v]$, respectively. Note that $x_{u,v} > 0$ for all edges $\{u, v\} \in E$ with $u \in V_{\max}$ and $v \notin V_{\max}$, and $x_{u,v} < 0$ for all $\{u, v\} \in E$ with $u \in V_{\min}$ and $v \notin V_{\min}$. Furthermore, we claim that

$$\Delta\pi_u < \Delta\pi_v \quad \text{for all } u \in V_{\max}, v \in V_{\min}. \tag{13}$$

Assume by contradiction that $\Delta\pi_u \geq \Delta\pi_v$ for $u \in V_{\max}$, $v \in V_{\min}$. Since $\text{str}(\pi) = \pi_u - \pi_v$, we get

$$\text{str}(d') = \text{str}(\pi') \geq \pi'_u - \pi'_v = \text{str}(\pi) + \Delta\pi_u - \Delta\pi_v \geq \text{str}(\pi) = \text{str}(d),$$

a contradiction.

Let $c := \min_{v \in V_{\min}} \Delta\pi_v$,

$$V_{\Delta}^+ := \{v \in V : \Delta\pi_v \geq c\}, \quad \text{and} \quad V_{\Delta}^- := V \setminus V_{\Delta}^+.$$

Due to (12), Lemma 12 implies that both $G[V_{\Delta}^+]$ and $G[V_{\Delta}^-]$ are connected. By definition we have $V_{\min} \subseteq V_{\Delta}^+$, and (13) implies $V_{\max} \subseteq V_{\Delta}^-$. In particular, it follows that $V_{\max} \cap V_{\Delta}^+ = \emptyset$. Let \bar{U} be the set of nodes containing V_{Δ}^+ and all nodes that can be reached from V_{Δ}^+ without visiting any node of V_{\max} . Let $U := V \setminus \bar{U}$. Then we have $V_{\max} \subseteq U \subseteq V_{\Delta}^-$; see Figure 5.

Claim 1. $G[U]$ is connected.

of the claim. Let $u, v \in U$, $P \in \mathcal{P}_{u,v}$, and assume that P contains some node in \bar{U} . Let $\{\alpha, \bar{\alpha}\}$ and $\{\omega, \bar{\omega}\}$, with $\alpha, \omega \in U$ and $\bar{\alpha}, \bar{\omega} \in \bar{U}$, be the first and last edge of P in the cut $[U, \bar{U}]$ between U and \bar{U} . Then, $\alpha, \omega \in V_{\max}$ since otherwise α or ω could be reached from \bar{U} without passing any node of V_{\max} , contradicting $\alpha, \omega \notin \bar{U}$ and the definition of \bar{U} . But since $G[V_{\max}]$ is connected, there is a path from u to v via α and ω by only using nodes within U . \square

Similarly, let \bar{U}' be the set of nodes which contains V_{Δ}^- and all nodes that can be reached from V_{Δ}^- without visiting any node of V_{\min} , and define $U' := V \setminus \bar{U}'$. Then, $V_{\min} \subseteq U' \subseteq V_{\Delta}^+$. By the same line of arguments as before, $G[U']$ is connected. Note




Figure 5: Illustration of the different node sets in the proof of Theorem 14.

that, for every $\{u, v\} \in E$ with $u \in U$ and $v \in \bar{U}$, we have that $u \in V_{\max}$ and $v \in V \setminus V_{\max}$, and thus $x_{u,v} > 0$. Similarly, for every edge $\{u, v\} \in E$ with $u \in U'$ and $v \in \bar{U}'$, we have $x_{u,v} < 0$. Hence,

$$\sum_{v \in U} \text{bal}(x)_v > 0 \quad \text{and} \quad \sum_{v \in U'} \text{bal}(x)_v < 0. \quad (14)$$

As a consequence, there exists a source $s_i \in U \subseteq V_{\Delta}^-$ with $t_i \notin U$. Similarly, there exists a sink $t_k \in U' \subseteq V_{\Delta}^+$ with $s_k \notin U'$.

In the following, assume that (G, D) does not contain a type-1 minor. Then, due to (12), Lemma 13 implies

$$\Delta\pi_{s_\ell} \geq \Delta\pi_{t_\ell} \quad \text{for all } \ell \in I. \quad (15)$$

Since $s_i \in V_{\Delta}^-$ and $t_k \in V_{\Delta}^+$, (15) implies that $t_i \in V_{\Delta}^-$ and $s_k \in V_{\Delta}^+$. Therefore, (14) implies that $\bar{U} \subsetneq V_{\Delta}^-$ and $\bar{U}' \subsetneq V_{\Delta}^+$. Let W_1, \dots, W_q be the node sets of the connected components of the graph $G[V \setminus (U \cup U')]$. From the connectedness of $G[V_{\Delta}^+]$ and $G[V_{\Delta}^-]$ it follows that $[U, W_r] \neq \emptyset$ and $[U', W_r] \neq \emptyset$ for all $r = 1, \dots, q$; see Figure 5.

Claim 2. The graph $G[V \setminus (U \cup U')]$ is connected, that is, $q = 1$.

of the claim. Assume by contradiction that $q \geq 2$. Let R be a subset of $\{1, \dots, q\}$ of maximal cardinality such that, for every $r \in R$, there exists an $\ell \in I$ with $s_\ell \in U$ and $t_\ell \in W_r$. We claim that for every $r \in R$ it holds that

$$s_{\ell'} \notin W_r \quad \text{for all } \ell' \in I \text{ with } t_{\ell'} \notin W_r. \quad (16)$$

Otherwise, there exists $\ell' \in I$ with $t_{\ell'} \notin W_r$ and $s_{\ell'} \in W_r$ and, since $r \in R$ and by definition of R , there exists $\ell \in I$ with $s_\ell \in U$ and $t_\ell \in W_r$. But since $q \geq 2$, the subgraph $G[V \setminus W_r]$ is connected, and hence there exist node-disjoint paths $P \in \mathcal{P}_{s_\ell, t_{\ell'}}$ and $P' \in \mathcal{P}_{s_{\ell'}, t_\ell}$. Thus, by Lemma 11, (G, D) contains a type-1 minor, a contradiction.

Let $W_R := \bigcup_{r \in R} W_r$ and $U_R := W_R \cup U$. Then, since $x_{u,v} > 0$ for every $\{u, v\} \in E$ with $u \in U_R$ and $v \notin U_R$, we have

$$\sum_{v \in U_R} \text{bal}(x)_v > 0.$$

Therefore, there exists a source $s_\ell \in U_R$ with $t_\ell \notin U_R$; in particular, $t_\ell \notin W_R$. Thus, (16) implies $s_\ell \notin W_R$, and hence $s_\ell \in U$. Due to the maximality of R , $t_\ell \notin W_r$ for all $r \in \{1, \dots, q\}$, and hence $t_\ell \in U'$. Therefore, $\Delta\pi_{t_\ell} > \Delta\pi_{s_\ell}$, a contradiction to (15). \square




Figure 6: Construction of a type-2 network in the proof of Theorem 14.

Recall that there exists a source $s_i \in U \subset V_\Delta^-$ with $t_i \in V_\Delta^- \setminus U$ and a sink $t_k \in U' \subset V_\Delta^+$ with $s_k \in V_\Delta^+ \setminus U'$. By the above claim it follows that $t_i \in V_\Delta^- \cap W_1$ and $s_k \in V_\Delta^+ \cap W_1$. Moreover, due to (12), the fact that Δx is an s_j - t_j -flow, and the definitions of V_Δ^+ and V_Δ^- , Lemma 12 implies that $s_j \in V_\Delta^+$ and $t_j \in V_\Delta^-$.

We use our insights to construct a type-2 minor; see Figure 6. Let P be an s_k - t_i -path which neither contains s_i nor t_k (such a path exists since $s_k, t_i \in W_1$ and $s_i, t_k \notin W_1$) and which uses a minimal number of edges of the cut $[V_\Delta^+, V_\Delta^-]$. Let $P_+ := P \cap G[V_\Delta^+]$ and $P_- := P \cap G[V_\Delta^-]$ be the parts of P which are contained in $G[V_\Delta^+]$ and $G[V_\Delta^-]$, respectively. Furthermore, let P_+^1, \dots, P_+^m and P_-^1, \dots, P_-^m be the connected components of P_+ and P_- , where P_+^1 contains s_k and P_-^m contains t_i ; see Figure 6.

Due to the connectedness of $G[V_\Delta^+]$, there exists a path from t_k to P_+^r for every $r \in \{1, \dots, m\}$. By the minimality of P , for every $r, r' \in \{1, \dots, m\}$ with $r \neq r'$, any path connecting P_+^r and $P_+^{r'}$ in $G[V_\Delta^+]$ contains t_k .

Let $P_{t_k} \subset G[V_\Delta^+]$ be a path from t_k to P_+^1 , and let $P_{s_j} \subset G[V_\Delta^+]$ be a path from s_j to $P_+ \cup P_{t_k}$. Then P_{s_j} ends in P_+^1 since otherwise there exist node-disjoint s_j - t_k - and s_k - t_j -paths, which would imply the existence of a type-1 minor. Let $P_{s_i} \subset G[V_\Delta^-]$ be a path from s_i to P_-^m , and let $P_{t_j} \subset G[V_\Delta^-]$ be a path from t_j to $P_- \cup P_{s_i}$. By a similar argument as before, it follows that P_{t_j} ends in P_-^m .

Let $e_i \in E(P_{s_i})$, $e_j \in E(P) \cap [V_\Delta^+, V_\Delta^-]$, and $e_k \in E(P_{t_k})$. Then, deleting all labels except for $s_i, t_i, s_j, t_j, s_k, t_k$, deleting all edges in E except for the ones in the tree $T := P \cup P_{t_k} \cup P_{s_j} \cup P_{s_i} \cup P_{t_j}$, and afterwards contracting all edges in $E(T) \setminus \{e_i, e_j, e_k\}$ yields a type-2 minor. This completes the proof. \square

As mentioned in Remark 2, as a consequence of Theorem 14, robustness and strong robustness are in fact equivalent.

Corollary 15. *A network topology is robust if and only if it is strongly robust.*

Proof. By Theorem 14, a network topology is robust if and only if it neither contains a type-1 nor a type-2 minor. This condition is independent of the choice of the functions $\psi_{u,v}, \{u, v\} \in E$. \square

Furthermore, we can conclude that every network topology containing only a single source or a single sink is robust.

Corollary 16. *Let (G, D) be a network topology with a single source or a single sink, i.e., $|\bigcup_{i \in I} \{s_i\}| = 1$ or $|\bigcup_{i \in I} \{t_i\}| = 1$. Then (G, D) is robust.*

Proof. Every type-1 and every type-2 network contains at least two distinct sources and two distinct sinks. Thus, (G, D) does not contain a type-1 or type-2 network as a minor. By Theorem 14, (G, D) is robust. \square

To conclude this section, we prove two lemmas which turn out to be useful later. The first one states that in a network topology, edges which are not contained in any s_i - t_i -path never carry any flow. Based on this fact, the second lemma concludes that contracting these edges has no influence on the robustness of the network topology.

Lemma 17. *Let (G, D) be a network topology, $d \in \mathbb{R}_{\geq 0}^I$, $\pi = f^{-1}(\text{bal}(d))$, and $x = x(\pi)$ the flow satisfying demands d . If $x_{u,v} \neq 0$ for some $e = \{u, v\} \in E$, then e is contained in an s_i - t_i -path for some $i \in I$.*

Proof. We proceed by induction on $p(d) := |\{i \in I : d_i > 0\}|$. If $p(d) = 0$, then $x = 0$ and there is nothing to show. Now assume $p(d) > 0$, and consider $\{u, v\} \in E$ with $x_{u,v} \neq 0$. Let $j \in I$ with $d_j > 0$ and define $d' \in \mathbb{R}_{\geq 0}^I$ by $d'_j := 0$ and $d'_i := d_i$ for all $i \in I \setminus \{j\}$. Let $\pi' := f^{-1}(\text{bal}(d'))$ and $x' := x(\pi')$ the flow satisfying demands d' . Note that $p(d') = p(d) - 1$. If $x'_{u,v} \neq 0$, then, by induction, edge $\{u, v\}$ is contained in an s_i - t_i -path for some $i \in I$. Thus, we can assume that $x'_{u,v} = 0$. Let $\Delta x := x - x'$ and $\Delta \pi := \pi - \pi'$. Then Δx is an s_j - t_j -flow of value $d_j > 0$ and, due to Equation (1) and Assumption 1, $\text{sgn}(\Delta x_{u,v}) = \text{sgn}(\Delta \pi_u - \Delta \pi_v)$ for all $\{u, v\} \in E$. Furthermore, since $x_{u,v} \neq 0$ and $x'_{u,v} = 0$, we have $\Delta x_{u,v} \neq 0$. Hence, by Lemma 12(b) we conclude that edge $\{u, v\}$ is contained in some s_j - t_j -path. \square

Lemma 18. *Let (G, D) be a network topology and (\bar{G}, \bar{D}) its minor obtained by contracting all edges $e \in E$ that are not contained in any s_i - t_i -path, $i \in I$. Then (G, D) is robust if and only if (\bar{G}, \bar{D}) is robust.*

Proof. “ \Rightarrow ”: This direction follows directly from Lemma 7.

“ \Leftarrow ”: Let $d \in \mathbb{R}_{\geq 0}^I$, $\pi := f^{-1}(\text{bal}(d))$, and $x := x(\pi)$ be the flow satisfying demands d . By Lemma 17, $x_{u,v} = 0$ and thus $\pi_u = \pi_v$ for all edges $e = \{u, v\} \in E$ which are not contained in any s_i - t_i -path, $i \in I$. Therefore, contracting these edges does not alter the stress on the network, i.e., $\text{str}_G(d) = \text{str}_{\bar{G}}(d)$. Thus, due to the robustness of (\bar{G}, \bar{D}) , for every $d' \in \mathbb{R}_{\geq 0}^I$ with $d' \leq d$, we have $\text{str}_G(d') = \text{str}_{\bar{G}}(d') \leq \text{str}_{\bar{G}}(d) = \text{str}_G(d)$. Hence (G, D) is robust. \square

4 Robustness for special graph classes

For certain special classes of graphs, such as trees and cycles, we can give more explicit characterizations of robustness.

4.1 Robust tree topologies

For the case that the network is a tree, it turns out that robustness is closely related to the possibility to give an orientation to each edge of the tree such that the resulting directed graph is a so-called *bi-arborescence* (see Figure 7) and the unique paths from a source to the corresponding sink follow this orientation.

We use the following terminology: A directed graph is an *arc-tree* if its underlying undirected graph is a tree. Furthermore, for a directed graph (V, A) , a sequence of pairwise distinct nodes $P = (v_1, \dots, v_k)$ is a *path* if $(v_i, v_{i+1}) \in A$ or $(v_{i+1}, v_i) \in A$ for all $i = 1, \dots, k-1$, and P is a *directed path* if $(v_i, v_{i+1}) \in A$ for all $i = 1, \dots, k-1$. Denote the arc set of P by $A(P)$. For two nodes $u, v \in V$, a (directed) path (v_1, \dots, v_k) is a (directed) *u-v-path* if $v_1 = u$ and $v_k = v$.




Figure 7: Three examples of bi-arborescences.

Definition 19. An arc-tree (V, A) is called a *bi-arborescence* if there exists a node $r \in V$, called the *root*, such that, for every $v \in V$, there exists a directed v - r -path or a directed r - v -path.

See Figure 7 for an illustration. Note that any arborescence is also a bi-arborescence. In contrast to arborescences, however, the root of a bi-arborescence is not necessarily unique.

We give an alternative characterization of bi-arborescences in terms of changes of edge orientations along a path. To that end, for a directed graph (V, A) and a node $v \in V$, let $\delta^+(v) \subseteq A$ be the set of arcs that start in v , and let $\delta^-(v) \subseteq A$ be the set of arcs that end in v .

Definition 20. Let (V, A) be a directed graph and P a path. A node $v \in V$ with $|\delta^-(v) \cap A(P)| = 2$ or $|\delta^+(v) \cap A(P)| = 2$ is called a *flipping node* of P . Moreover, the number

$$\varphi(P) := |\{v \in V : v \text{ is a flipping node of } P\}|$$

is called the *number of flips* of path P .

Lemma 21. An arc-tree (V, A) is a bi-arborescence if and only if every path in (V, A) has at most one flip.

Proof. “ \Rightarrow ”: If (V, A) is a bi-arborescence, then, for any path P , the only possible flipping node of P is the node with minimal graph-theoretic distance to the root r .

“ \Leftarrow ”: If $|\delta^-(v)| \leq 1$ for all $v \in V$, then (V, A) is an arborescence and thus a bi-arborescence. Otherwise, among all nodes u with $|\delta^-(u)| \geq 2$ choose one such that there is no directed path from u to any other such node (this node exists due to the acyclicity of the graph). We call this node r and argue that (V, A) is a bi-arborescence with root r : For any node $v \in V$ such that the first arc on the unique r - v -path is in $\delta^+(r)$, this path is directed by our choice of r . Moreover, for any node $v \in V$ such that the first arc on the unique r - v -path is in $\delta^-(r)$, the reverse path is a directed v - r -path, since otherwise adding the other arc in $\delta^-(r)$ to the r - v -path yields a path with at least two flips, a contradiction. \square

We proceed to give a characterization of robust trees. It will depend on the possibility to turn the tree into a bi-arborescence such that each s_i - t_i -path follows this orientation, i.e., each s_i - t_i -path is directed.




Figure 8: An arc-tree containing a type-2 minor, as described in the proof of Theorem 24.

Definition 22. For a simple undirected graph (V, E) , a set of arcs $A \subset V \times V$ is called an *orientation* of E if $|A| = |E|$ and $E = \{\{u, v\} : (u, v) \in A\}$. A network topology (T, D) is called *tree topology* if $T = (V, E)$ is a tree. A tree topology is *regular* if there exists an orientation A of E such that (V, A) is a bi-arborescence with a directed \mathfrak{s}_i - \mathfrak{t}_i -path for every $i \in I$.

Lemma 23. *Every minor of a regular tree topology is regular.*

Proof. Let (T, D) be a regular tree topology. Note that contracting an edge $e \in E$ or deleting a pair of labels $\mathfrak{s}_i, \mathfrak{t}_i$ does not destroy the regularity of the network topology. Since T is a tree, every minor of (T, D) can be obtained by a finite sequence of edge contractions and label pair deletions. By induction, every minor of (T, D) is regular. \square

For a tree topology (T, D) denote by P_i^T the unique \mathfrak{s}_i - \mathfrak{t}_i -path in T , $i \in I$. We can now give a characterization of robust tree topologies.

Theorem 24. *Let (T, D) be a tree topology and let (\bar{T}, \bar{D}) be the minor that is obtained from (T, D) by contracting all edges $e \in E$ with $e \notin \bigcup_{i \in I} E(P_i^T)$. Then (T, D) is robust if and only if (\bar{T}, \bar{D}) is regular.*

Proof. By Lemma 18, (T, D) is robust if and only if (\bar{T}, \bar{D}) is robust. Thus, it remains to show that (\bar{T}, \bar{D}) is robust if and only if it is regular.

“ \Rightarrow ”: We show the contraposition. First, assume that there is no orientation \bar{A} of \bar{E} such that the arc-tree (\bar{V}, \bar{A}) contains a directed \mathfrak{s}_i - \mathfrak{t}_i -path for every $i \in I$. Then, there exist $i, j \in I$ and an edge $e \in E(P_i^T) \cap E(P_j^T)$ such that e is traversed in opposite directions when going from \mathfrak{s}_i to \mathfrak{t}_i and from \mathfrak{s}_j to \mathfrak{t}_j , respectively. Thus, contracting every edge in $\bar{E} \setminus \{e\}$ and deleting all labels except for $\mathfrak{s}_i, \mathfrak{t}_i, \mathfrak{s}_j, \mathfrak{t}_j$ yields a type-1 minor. By Theorem 14, (\bar{T}, \bar{D}) is not robust.

Second, assume there is an orientation \bar{A} of \bar{E} such that, for every $i \in I$, the arc-tree (\bar{V}, \bar{A}) contains a directed \mathfrak{s}_i - \mathfrak{t}_i -path, but (\bar{V}, \bar{A}) is not a bi-arborescence. By Lemma 21, (\bar{V}, \bar{A}) contains a path P with two flipping nodes. Thus, as depicted in Figure 8, there are three arcs $a_i, a_j, a_k \in A(P)$ such that, when traversing P ,

- a_i is traversed before a_j ,
- a_j is traversed before a_k ,
- a_i and a_k are traversed along their orientation,
- a_j is traversed against its orientation.

Furthermore, by the construction of the arc-tree (\bar{V}, \bar{A}) , there exist $i, j, k \in I$ such that, for every $\ell \in \{i, j, k\}$, the arc a_ℓ is contained in the directed \mathfrak{s}_ℓ - \mathfrak{t}_ℓ -path. Let $e_i, e_j, e_k \in \bar{E}$ be the edges corresponding to the arcs $a_i, a_j, a_k \in \bar{A}$. Within (\bar{T}, \bar{D}) , contracting all edges in $\bar{E} \setminus \{e_i, e_j, e_k\}$ and deleting all labels except for $\mathfrak{s}_i, \mathfrak{t}_i, \mathfrak{s}_j, \mathfrak{t}_j, \mathfrak{s}_k, \mathfrak{t}_k$ yields a type-2 minor. By Theorem 14, (\bar{T}, \bar{D}) is not robust.

“ \Leftarrow ”: Assume (\bar{T}, \bar{D}) is regular. Then, by Lemma 23, every minor of (\bar{T}, \bar{D}) is regular. But then (\bar{T}, \bar{D}) can neither contain a type-1 nor a type-2 minor, since these minors are clearly not regular. By Theorem 14, (\bar{T}, \bar{D}) is robust. \square




Figure 9: The cycles in (a) and (b) are robust, whereas the cycles in (c) and (d) are not robust. The cycle in (c) violates Property (a) in Theorem 26. The cycle in (d) violates Property (b) since $s_3 \prec_{P_u} s_2$ and $t_2 \prec_{P_v} t_3$.

4.2 Robust cycles

Next, we give a characterization of robust cycles. We start by considering a necessary condition for robustness.

Lemma 25. *Let (G, D) be a network topology such that G is a cycle. If (G, D) is robust, then, for all $i, j \in I$, every s_i - t_j -path contains s_j or t_i .*

Proof. If there exists an s_i - t_j -path which neither contains s_j nor t_i , then there also exists an s_j - t_i -path which neither contains s_i nor t_j . Thus, by Lemma 11, (G, D) contains a type-1 minor and is therefore not robust. \square

We can now give a full characterization of robust cycles. To that end, for a path $P = (v_1, \dots, v_k)$ we write $v_i \prec_P v_j$ if $i < j$.

Theorem 26. *Let (G, D) be a network topology such that G is a cycle and $s_i \neq t_i$ for all $i \in I$. Then (G, D) is robust if and only if both of the following two conditions hold:*

- (a) *There exist two edges $\{u_1, v_1\}, \{u_2, v_2\} \in E$ and two node-disjoint paths $P_u \in \mathcal{P}_{u_1, u_2}$, $P_v \in \mathcal{P}_{v_1, v_2}$ with $\{s_i : i \in I\} \subseteq V(P_u)$, $\{t_i : i \in I\} \subseteq V(P_v)$.*
- (b) *Moreover, there is no pair $i, j \in I$ such that $s_i \prec_{P_u} s_j$ and $t_j \prec_{P_v} t_i$.*

Figure 9 shows examples of robust and non-robust cycles.

of Theorem 26. “ \Rightarrow ”: By contradiction, first assume that (a) holds, but (b) does not hold. Then there exist $i, j \in I$ with $s_i \prec_{P_u} s_j$ and $t_j \prec_{P_v} t_i$. Consequently, there exist node-disjoint paths $P \in \mathcal{P}_{s_i, t_j}$ and $P' \in \mathcal{P}_{s_j, t_i}$. By Lemma 11, (G, D) contains a type-1 minor and is thus not robust, a contradiction.

Next, assume that (a) does not hold. Thus, there exist $i, j, k \in I$ such that the cycle G can be decomposed into the concatenation of four paths: $P_1 \in \mathcal{P}_{s_i, t_j}$, $P_2 \in \mathcal{P}_{t_j, s_k}$, $P_3 \in \mathcal{P}_{s_k, t_i}$, and $P_4 \in \mathcal{P}_{t_i, s_i}$; see Figure 10.

Applying Lemma 25 to s_i and t_j implies that $s_j \in V(P_1)$. Applying Lemma 25 once more to s_k and t_j implies that $t_k \in V(P_2)$. But then there exists an s_i - t_k -path which neither contains s_k nor t_i , contradicting Lemma 25.




Figure 10: A cycle violating Property (a) of Theorem 26.




Figure 11: The possible cases for non-robust cycles. In (a) the cycle contains a type-2 minor. In (b) and (c) it contains a type-1 minor.

“ \Leftarrow ”: Assume by contradiction that (G, D) is not robust. By Theorem 14, it contains a type-1 or a type-2 minor. First assume that (G, D) contains a type-2 minor. Then G is of the form shown in Figure 11a where each of the paths P_2 and P_4 might be of length zero, whereas all of the paths $P_1, P_3, P_5,$ and P_6 contain at least one edge. Consequently, Property (a) is not satisfied. Now, assume (G, D) contains a type-1 minor. Then G has one of the forms shown in Figure 11b and Figure 11c, where each of the paths P_1 and P_3 might be of length zero, whereas both of the paths P_2 and P_4 contain at least one edge. In the first case, Property (a) is violated. In the second case, if P_1 or P_3 has length zero, then Property (a) is violated, otherwise Property (b) is violated. \square

5 Robustness in the entry-exit model

Let $G = (V, E)$ be an undirected graph, and let $S, T \subset V$ be nonempty disjoint sets of sources and sinks, respectively. We call (G, S, T) an *entry-exit topology*. In this section we define robustness for entry-exit topologies and give a complete characterization of the class of robust entry-exit topologies by exploiting the results on robustness of network topologies proved in the previous sections.

The set of possible balance vectors is

$$B(S, T) := \left\{ b \in \mathbb{R}^V : b_s \geq 0 \text{ for all } s \in S, b_t \leq 0 \text{ for all } t \in T, \sum_{v \in V} b_v = 0, b_v = 0 \text{ for all } v \in V \setminus (S \cup T) \right\}.$$

For $b, b' \in B(S, T)$ we write $b \preceq b'$ if $|b_v| \leq |b'_v|$ for all $v \in V$.

Definition 27. An entry-exit topology (G, S, T) together with potential loss functions $\psi_{u,v} \in \Psi$, for $\{u, v\} \in E$, is called *robust* if, for all $\beta \in \mathbb{R}_{>0}^E$, the function $\text{str}: B(S, T) \rightarrow \mathbb{R}_{\geq 0}$ is non-decreasing with respect to \preceq , i.e., for all $b, b' \in B(S, T)$ with $b \preceq b'$ we have $\text{str}(b) \leq \text{str}(b')$.

We will show that robustness of an entry-exit topology is equivalent to robustness of a certain corresponding network topology. To that end, let

$$D(S, T) := ((s, t) : (s, t) \in S \times T)$$

be the tuple of all source-sink pairs with sources in S and sinks in T . Sticking to the notation of the previous sections, this means that $D(S, T) = (D_i)_{i \in I}$, where $I = S \times T$, and for $i \in I$ we have $D_i = (s_i, t_i)$. We state the following classic observation:

Observation 1. *For each $b \in B(S, T)$ there is a $d \in \mathbb{R}_{\geq 0}^{S \times T}$ with $\text{bal}(d) = b$.*

Lemma 28. *For every $b, b' \in B(S, T)$ with $b \preceq b'$, there exist $d, d' \in \mathbb{R}_{\geq 0}^{S \times T}$ with $d \leq d'$ such that $\text{bal}(d) = b$ and $\text{bal}(d') = b'$.*

Proof. Let $b, b' \in B(S, T)$ with $b \preceq b'$ and set $\Delta b := b' - b \in B(S, T)$. By Observation 1, let $d, \Delta d \in \mathbb{R}_{\geq 0}^{S \times T}$ with $\text{bal}(d) = b$ and $\text{bal}(\Delta d) = \Delta b$, and set $d' := d + \Delta d \in \mathbb{R}_{\geq 0}^{S \times T}$. Then $d \leq d'$ and $\text{bal}(d') = \text{bal}(d) + \text{bal}(\Delta d) = b + \Delta b = b'$. \square

Lemma 29. *An entry-exit topology (G, S, T) is robust if and only if the network topology $(G, D(S, T))$ is robust in the sense of Definition 1.*

Proof. “ \Rightarrow ”: Assume that (G, S, T) is robust, and let $d \leq d' \in \mathbb{R}_{\geq 0}^{S \times T}$. Then,

$$\begin{aligned} \text{bal}(d)_s &= \sum_{t \in T} d_{(s,t)} \leq \sum_{t \in T} d'_{(s,t)} = \text{bal}(d')_s \quad \text{for all } s \in S, \\ \text{bal}(d)_t &= - \sum_{s \in S} d_{(s,t)} \geq - \sum_{s \in S} d'_{(s,t)} = \text{bal}(d')_t \quad \text{for all } t \in T, \\ \text{bal}(d)_v &= 0 = \text{bal}(d')_v \quad \text{for all } v \in V \setminus (S \cup T). \end{aligned}$$

Thus, $\text{bal}(d) \preceq \text{bal}(d')$ and $\text{bal}(d), \text{bal}(d') \in B(S, T)$. Since (G, S, T) is robust, it follows that $\text{str}(d) = \text{str}(\text{bal}(d)) \leq \text{str}(\text{bal}(d')) = \text{str}(d)$, showing that $(G, D(S, T))$ is robust.

“ \Leftarrow ”: Assume that $(G, D(S, T))$ is robust, and let $b, b' \in B(S, T)$ with $b \preceq b'$. By Lemma 28 there exist $d, d' \in \mathbb{R}_{\geq 0}^{S \times T}$ with $d \leq d'$, $\text{bal}(d) = b$, and $\text{bal}(d') = b'$. Since $(G, D(S, T))$ is robust, it follows that $\text{str}(b) = \text{str}(\text{bal}(d)) \leq \text{str}(\text{bal}(d')) = \text{str}(b')$, showing that (G, S, T) is robust. \square

As our main result of this section, we provide a complete characterization of robustness for entry-exit topologies. The following theorem states that an entry-exit topology is robust if and only if there is an articulation node separating sources from sinks, i.e., there exists a node $r \in V$ such that removing r from G disconnects G into two subgraphs G_S and G_T with $S \setminus \{r\} \subseteq V(G_S)$ and $T \setminus \{r\} \subseteq V(G_T)$.

Theorem 30. *An entry-exit topology (G, S, T) is robust if and only if G contains an articulation node which separates the sources from the sinks.*

Proof. By Lemma 29 it suffices to show that the network topology $(G, D(S, T))$ is robust if and only if G contains an articulation node that separates the sources from the sinks.

“ \Rightarrow ”: Assume that there is no articulation node separating the sources from the sinks. Applying Menger’s Theorem to S and T yields the existence of $s, s' \in S, t, t' \in T$ such that there are two node-disjoint paths $P \in \mathcal{P}_{s,t}$ and $P' \in \mathcal{P}_{s',t'}$. For $i := (s, t') \in I$ and $j := (s', t) \in I$, we have that $\mathcal{P}_{s,t} = \mathcal{P}_{s_i,t_j}$ and $\mathcal{P}_{s',t'} = \mathcal{P}_{s_j,t_i}$. Consequently, Lemma 11 implies that $(G, D(S, T))$ contains a type-1 minor. By Theorem 14, $(G, D(S, T))$ is not robust.

“ \Leftarrow ”: Assume that $(G, D(S, T))$ is not robust. Then Theorem 14 implies that $(G, D(S, T))$ contains a type-1 or a type-2 minor. First, assume that $(G, D(S, T))$ contains a type-1 minor. By Lemma 11, there exists $i = (s, t') \in I$ and $j = (s', t) \in I$ and node-disjoint paths $P \in \mathcal{P}_{s_i, t_j} = \mathcal{P}_{s, t}$ and $P' \in \mathcal{P}_{s_j, t_i} = \mathcal{P}_{s', t'}$. Consequently, there exists no articulation node between the sources and the sinks.

Next, assume that $(G, D(S, T))$ contains a type-2 network (\bar{G}, \bar{D}) as a minor; see Figure 1b. This type-2 minor contains node disjoint paths $\bar{P} \in \mathcal{P}_{\bar{s}_i, \bar{t}_i}^{\bar{G}}$ and $\bar{P}' \in \mathcal{P}_{\bar{s}_k, \bar{t}_k}^{\bar{G}}$. By Lemma 10, also $(G, D(S, T))$ contains node-disjoint paths $P \in \mathcal{P}_{s_i, t_i}^G$ and $P' \in \mathcal{P}_{s_k, t_k}^G$ and can therefore not contain an articulation node separating the sources from the sinks. \square

6 Algorithmically determining the robustness of a network

After having fully characterized robust network topologies theoretically, we show in this section that the robustness of a network topology can be decided in polynomial time. Referring to Theorem 14, deciding the robustness of a network is equivalent to deciding whether a network has a type-1 or type-2 network as a minor. In this context, Robertson and Seymour [30] showed that, for a fixed graph \bar{G} , there is an algorithm which decides whether \bar{G} is a minor of larger graph G in time polynomial in the size of G . Their result, however, is only known to hold for unlabeled graph minors, and a generalization of their result to labeled graph minors is beyond the scope of this work. Instead, to show that we can efficiently decide whether a graph has a type-1 or type-2 minor, we will use the following result of Robertson and Seymour [30].

Lemma 31 (Robertson and Seymour [30]). *Let $G = (V, E)$ be an undirected graph, k a fixed natural number, and $\mathcal{A} = \{(u_1, v_1), \dots, (u_k, v_k)\} \subseteq V \times V$. Then, the existence of node-disjoint paths $P_1 \in \mathcal{P}_{u_1, v_1}, \dots, P_k \in \mathcal{P}_{u_k, v_k}$ can be decided in time polynomial in the size of G .*

For a set of node pairs $\mathcal{A} = \{(u_1, v_1), \dots, (u_k, v_k)\} \subseteq V \times V$, we call paths $P_1 \in \mathcal{P}_{u_1, v_1}, \dots, P_k \in \mathcal{P}_{u_k, v_k}$ *internally node-disjoint* if, for all $i, j \in \{1, \dots, k\}$ with $i \neq j$, we have $(V(P_i) \setminus \{u_i, v_i\}) \cap V(P_j) = \emptyset$, i.e., all nodes in $V(P_i)$ and $V(P_j)$ are pairwise distinct, except for possibly $\{u_i, v_i\} \cap \{u_j, v_j\} \neq \emptyset$. Note that deciding whether a graph G contains *internally* node-disjoint paths $P_i \in \mathcal{P}_{u_i, v_i}^G$, for $i = 1, \dots, k$, can be easily reduced in polynomial time to deciding whether a certain slightly modified graph \hat{G} contains node-disjoint paths $\bar{P}_i \in \mathcal{P}_{\hat{u}_i, \hat{v}_i}^{\hat{G}}$, where (\hat{u}_i, \hat{v}_i) are corresponding copies of the pairs (u_i, v_i) , $i = 1, \dots, k$. Let $(\text{INTERNALLY})\text{NODEDISJOINTPATHS}(\mathcal{A})$ be the polynomial algorithm that decides whether there exist (internally) node-disjoint paths $P_i \in \mathcal{P}_{u_i, v_i}$, for $i = 1, \dots, k$.

Theorem 32. *Deciding whether a network topology (G, D) is robust can be done in time polynomial in the input size of (G, D) .*

Proof. By Lemma 11, (G, D) contains a type-1 minor if and only if there exist $i, j \in I$ with two node-disjoint paths $P \in \mathcal{P}_{s_i, t_j}$ and $P' \in \mathcal{P}_{s_j, t_i}$. This can be decided in polynomial time by calling $\text{NODEDISJOINTPATHS}(\{(s_i, t_j), (s_j, t_i)\})$ for all $i, j \in I$ with $i \neq j$.

Moreover, if (G, D) does not contain a type-1 minor, we claim that the polynomial algorithm $\text{CONTAINSTYPE2MINOR}$ correctly decides whether (G, D) contains a type-2 minor. To prove this claim, assume that (G, D) contains a type-2 minor (\bar{G}, \bar{D}) , but no type-1 minor. Denote the three edges of \bar{G} by e_i, e_j , and e_k ; see Figure 12. By Lemma 4, G contains a subtree for which contracting all edges except e_i, e_j , and e_k , and deleting

Algorithm 1: CONTAINSTYPE2MINOR

Input: Network topology (G, D) without a type-1 minor
Output: Decision whether (G, D) contains a type-2 minor

```

1 for  $(s_i, t_i), (s_j, t_j), (s_k, t_k) \in D$  such that  $s_i, t_i, s_j, t_j, s_k, t_k$  are pairwise distinct
  except for possibly  $t_i = t_j$  and/or  $s_j = s_k$  do
2    $U \leftarrow V \setminus \{s_i, t_i, s_j, t_j, s_k, t_k\}$ 
3   for  $u_i, u_j \in U \cup \{t_i, t_j\}$  do
4     for  $v_j, v_k \in (U \setminus \{u_i, u_j\}) \cup \{s_j, s_k\}$  do
5       /* the different possibilities for the  $s_i$ - $t_k$ -path  $P$  */
6        $\mathcal{A}_1 \leftarrow \{(s_i, u_j), (u_j, v_j), (v_j, t_k)\}$ 
7        $\mathcal{A}_2 \leftarrow \{(s_i, u_i), (u_i, u_j), (u_j, v_j), (v_j, t_k)\}$ 
8        $\mathcal{A}_3 \leftarrow \{(s_i, u_j), (u_j, v_j), (v_j, v_k), (v_k, t_k)\}$ 
9        $\mathcal{A}_4 \leftarrow \{(s_i, u_i), (u_i, u_j), (u_j, v_j), (v_j, v_k), (v_k, t_k)\}$ 
10      /* connecting  $P$  to  $t_i$  and  $t_j$  */
11      for  $\ell = 1, 3$  do
12         $\mathcal{A}_\ell \leftarrow \mathcal{A}_\ell \cup \{(u_j, u_i), (u_i, t_i), (u_i, t_j)\}$ 
13      for  $\ell = 2, 4$  do
14         $\mathcal{A}_\ell \leftarrow \mathcal{A}_\ell \cup \{(u_i, t_i), (u_j, t_j)\}$ 
15      /* connecting  $P$  to  $s_j$  and  $s_k$  */
16      for  $\ell = 1, 2$  do
17         $\mathcal{A}_\ell \leftarrow \mathcal{A}_\ell \cup \{(v_j, v_k), (v_k, s_j), (v_k, s_k)\}$ 
18      for  $\ell = 3, 4$  do
19         $\mathcal{A}_\ell \leftarrow \mathcal{A}_\ell \cup \{(v_j, s_j), (v_k, s_k)\}$ 
20      for  $\ell = 1, 2, 3, 4$  do
21        if INTERNALLYDISJOINTPATHS( $\mathcal{A}_\ell$ ) then
22          return true
23 return false

```




Figure 12: The type-2 minor in the proof of Theorem 32.

all labels except for $s_i, t_i, s_j, t_j, s_k, t_k$ yields (\bar{G}, \bar{D}) . Within this tree, let P be the path from s_i to t_k , let P_{t_j} and P_{s_j} be the paths from t_j to P and from s_j to P , respectively, and let P_{t_i} and P_{s_k} be the paths from t_i to $P \cup P_{t_j}$ and from s_k to $P \cup P_{s_j}$, respectively. Finally, let $u_j \in V(P_{t_j}) \cap V(P)$, $v_j \in V(P_{s_j}) \cap V(P)$, $u_i \in V(P_{t_i}) \cap V(P \cup P_{t_j})$, and $v_k \in V(P_{s_k}) \cap V(P \cup P_{s_j})$; see Figure 13 for an illustration. Four cases can occur:

1. $u_i \notin V(P)$ and $v_k \notin V(P)$,
2. $u_i \in V(P)$ and $v_k \notin V(P)$,
3. $u_i \notin V(P)$ and $v_k \in V(P)$,
4. $u_i \in V(P)$ and $v_k \in V(P)$.

If $u_i \in V(P)$, then $u_i \preceq_P u_j$, since otherwise there exist node-disjoint s_i - t_j - and s_j - t_i paths, and hence (G, D) would contain a type-1 minor. Similarly, $v_k \in V(P)$ implies $v_j \preceq v_k$. Therefore, for the ℓ -th case, there exist internally node-disjoint paths for all pairs in \mathcal{A}_ℓ , $\ell \in \{1, 2, 3, 4\}$, as defined in lines 5–16 of Algorithm CONTAINSTYPE2MINOR. Thus, the algorithm returns **true**.

On the other hand, assume that the algorithm returns **true**, and let \mathcal{A}_ℓ , $\ell \in \{1, 2, 3, 4\}$, be the set of pairs that led to this output. Let T' be the union of the




Figure 13: Example of a tree T containing a type-2 minor. Here, $u_i \notin V(P)$ and $v_k \in V(P)$. Thus, T is the union of internally node-disjoint paths for the pairs in \mathcal{A}_3 .

internally node-disjoint paths between the pairs in \mathcal{A}_ℓ , and let $T \subseteq T'$ be a spanning tree of T' . Let P be the s_i - t_k -path in T , let e_i and e_k be the first and the last edge of P , respectively, and let e_j be an edge of the u_j - v_j -path in T . Then, deleting from G all edges in $E \setminus E(T)$, afterwards contracting all edges in $E(T) \setminus \{e_i, e_j, e_k\}$, and finally deleting all labels except for $s_i, t_i, s_j, t_j, s_k, t_k$ yields a type-2 minor. \square

References

- [1] Álvarez-Miranda, E., Cacchiani, V., Lodi, A., Parriani, T., Schmidt, D.R.: Single-commodity robust network design problem: Complexity, instances and heuristic solutions. *Eur. J. Oper. Res.* **238**, 711–723 (2014). DOI 10.1016/j.ejor.2014.04.023
- [2] Ben-Tal, A., Ghaoui, L.E., Nemirovski, A.S.: *Robust Optimization*. Princeton Series in Applied Mathematics. Princeton University Press (2009)
- [3] Ben-Tal, A., Nemirovski, A.: Robust solutions of uncertain linear programs. *Oper. Res. Lett.* **25**, 1–13 (1999). DOI 10.1016/S0167-6377(99)00016-4
- [4] Bienstock, O., Günlük, O.: Capacitated network design – Polyhedral structure and computation. *INFORMS J. Comput.* **8**, 243–259 (1996). DOI 10.1287/ijoc.8.3.243
- [5] Birkhoff, G., Diaz, J.B.: Non-linear network problems. *Quart. Appl. Math.* **13**(4), 431–443 (1956)
- [6] Braess, D.: Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung* **12**, 258–268 (1968)
- [7] Buchheim, C., Liers, F., Sanità, L.: An exact algorithm for robust network design. In: J. Pahl, T. Reiners, S., Voß (eds.) *Network Optimization. INOC 2011, Lecture Notes in Computer Science*, vol. 6701 (2011). DOI 10.1007/978-3-642-21527-8_2
- [8] Cacchiani, V., Jünger, M., Liers, F., Lodi, A., Schmidt, D.R.: Single-commodity robust network design with finite and Hose demand sets. *Math. Program., Ser. B* **157**, 297–342 (2016). DOI 10.1007/s10107-016-0991-9
- [9] Chien, R.T.: Synthesis of a communication net. *IBM J. Res. Develop.* **4**(3), 311–320 (1960). DOI 10.1147/rd.43.0311
- [10] Chou, W., Frank, H.: Survivable communication networks and the terminal capacity matrix. *IEEE Trans. on Circuit Theory* **17**, 192–197 (1970). DOI 10.1109/TCT.1970.1083100
- [11] Collins, M., Cooper, L., Helgason, R., Kennington, J., LeBlanc, L.: Solving the pipe network analysis problem using optimization techniques. *Management Sci.* **24**(7), 747–760 (1978). DOI 10.1287/mnsc.24.7.747
- [12] Diestel, R.: *Graph theory*. Springer-Verlag Berlin Heidelberg (2018)
- [13] Duffield, N.G., Goyal, P., Greenberg, A., Mishra, P., Ramakrishnan, K.K., van der Merwe, J.E.: A flexible model for resource management in virtual private networks. In: *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pp. 95–108 (1999). DOI 10.1145/316188.316209

- [14] Eisenbrand, F., Grandoni, F., Oriolo, G., Skutella, M.: New approaches for virtual private network design. *SIAM J. Comput* **37**, 706–721 (2007). DOI 10.1137/060654827
- [15] Fingerhut, J.A., Suri, S., Turner, J.S.: Designing least-cost nonblocking broadband networks. *J. Algorithms* **24**(2), 287–309 (1997). DOI 10.1006/jagm.1997.0866
- [16] Ford L.R., J., Fulkerson, D.R.: A suggested computation for maximal multi-commodity network flows. *Management Sci.* **5**, 97–101 (1958)
- [17] Friedmann, H., Robertson, N., Seymour, P.D.: The metamathematics of the graph minor theorem. *Contemporary Mathematics* **65**, 229–261 (1987)
- [18] Gomory, R.E., Hu, T.C.: Multi-terminal network flows. *J. Soc. Indust. Appl. Math.* **9**, 551–570 (1961). DOI 10.1137/0109047
- [19] Gomory, R.E., Hu, T.C.: An application of generalized linear programming to network flows. *J. Soc. Indust. Appl. Math.* **10**, 260–283 (1962)
- [20] Gomory, R.E., Hu, T.C.: Synthesis of a communication network. *J. Soc. Indust. Appl. Math.* **12**, 348–369 (1964). DOI 10.1137/0112029
- [21] Goyal, N., Olver, N., Shepherd, B.: The VPN conjecture is true. *J. ACM* **60**, 17:1–17:17 (2013). DOI 10.1145/2487241.2487243
- [22] Grimm, V., Schewe, L., Schmidt, M., Zöttl, G.: A multilevel model of the European entry-exit gas market. *Math. Methods Oper. Res.* (2018). DOI 10.1007/s00186-018-0647-z
- [23] Groß, M., Pfetsch, M.E., Schewe, L., Schmidt, M., Skutella, M.: Algorithmic results for potential-based flows: Easy and hard cases. *Networks* (2018). DOI 10.1002/net.21865
- [24] Gupta, A., Kleinberg, J., Kumar, A., Rastogi, R., Yener, B.: Provisioning a virtual private network: a network design problem for multicommodity flow. In: *Proc. 33rd Annu. ACM Sympos. Theory Comput. (STOC)*, pp. 389–398 (2001). DOI 10.1145/380752.380830
- [25] Gusfield, D.: Simple constructions for the multi-terminal network flow synthesis. *SIAM J. Comput* **12**, 157–165 (1983). DOI 10.1137/0212010
- [26] Kabadi, S.N., Yan, J., Du, D., Nair, K.P.K.: Integer exact network synthesis problem. *SIAM J. Discret. Math.* **23**(1), 136–154 (2008). DOI 10.1137/050641776
- [27] Koch, T., Hiller, B., Pfetsch, M.E., Schewe, L.: *Evaluating Gas Network Capacities*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA (2015). DOI 10.1137/1.9781611973693
- [28] Maugis, J.J.: Étude de réseaux de transport et de distribution de fluide. *RAIRO Oper. Res.* **11**(2), 243–248 (1977)
- [29] Minoux, M.: Optimum synthesis of a network with non-simultaneous multicommodity flow requirements. In: *Annals of Discrete Mathematics, Studies on Graphs and Discrete Programming*, vol. 59, pp. 269–277. North-Holland (1981). DOI 10.1016/S0304-0208(08)73470-4
- [30] Robertson, N., Seymour, P.D.: Graph minors. XIII. The disjoint paths problem. *J. Combin. Theory, Ser. B* **63**(1), 65–110 (1995). DOI 10.1006/jctb.1995.1006
- [31] Rockafellar, R.T.: *Network Flows and Monotropic Optimization*. Wiley-Interscience (1984)
- [32] Schmidt, D.R.: Robust design of single-commodity networks. Ph.D. thesis, Köln University (2014)
- [33] Sridhar, R., Chandrasekaran, R.: Integer solution to synthesis of communication networks. *Math. Oper. Res.* **17**, 581–585 (1992). DOI 10.1287/moor.17.3.581
- [34] Szabó, J.: The set of solutions to nomination validation in passive gas transportation networks with a generalized flow formula. Tech. Rep. 11-44, Zuse Institute Berlin (2012)
- [35] Talluri, K.T.: Network synthesis with few edges. *Networks* **27**, 109–115 (1996). DOI 10.1002/(SICI)1097-0037(199603)27:2<109::AID-NET2>3.0.CO;2-O
- [36] Weymouth, T.R.: Problems in natural gas engineering. *Transactions of the American Society of Mechanical Engineers* **34**, 185–231 (1912)