Directed Graph Embeddings

Reshef Meir
Technion—Israel Institute of Technology
reshefm@ie.technion.ac.il

November 7, 2017

Abstract

Definitions of graph embeddings and graph minors for directed graphs are introduced. For the class of 2-terminal directed acyclic graphs (TDAGs) the definitions coincide, and the class is closed under both operations. The usefulness of the directed embedding operation is demonstrated be characterizing all TDAGs with parallel-width at most k, which generalizes earlier characterizations of series-parallel graphs.

1 Introduction

Graph embeddings and graph minors have been used extensively in order to characterize classes of graphs with various properties (see e.g. (Kawarabayashi and Mohar, 2007) for a recent survey). Some famous results are the characterization of planar graphs (Kuratowski, 1930; Wagner, 1937), and of graphs with bounded treewidth (Robertson and Seymour, 1986). Similar results for directed graphs also exist (Kintali and Zhang, 2013; Johnson et al., 2001; Holzman and Law-Yone, 2003; Holzman and Monderer, 2015; Johnson et al., 2015) but are more scarce, and each one uses different graph operations.

We define new notions of graph embedding and graph minor for directed graphs. We show that for the class of 2-terminal directed acyclic graphs (TDAGs) these two operations exactly reverse one another, and thus a TDAG G' is a directed minor of G if and only if it is directed-embedded in G. Also, the class of TDAG is closed under directed minor and directed embedding operations.

Parallel width of a graph is a parameter recently introduced in the context of routing games (Meir and Parkes, 2017), and is useful for bounding negative externalities and the price of anarchy. For example, series-parallel graphs have parallel-width of 1, and the famous Braess paradox network has a parallel-width of 2, which intuitively means that there is a route intersecting two other disjoint routs (and thus a group of agents can negatively influence two other groups of agents).

For each k, we describe a finite set of graphs whose exclusion as directed minors of a TDAG G is necessary and sufficient to determine that G has parallel width at most k. For $k = 1$ there is only one forbidden directed minor, which is the Braess paradox graph. Thus our results extend known results for directed and undirected graphs (Holzman and Monderer, 2015; Milchtaich, 2006), albeit with somewhat different notions of graph embeddings. Some simple proofs are omitted and appear in the appendix.

For convenience, we will use the letter H for undirected graphs, and the letter G for directed graphs. We denote by $\overline{G}$ the undirected graph obtained from G by ignoring edge directions.

We denote a path in graph $\langle V, E\rangle$ by $(v_1, v_2, \dots, v_m)$, where for every i, $(v_i, v_{i+1}) \in E$. We use dash to abbreviate the path, e.g. $a - b - c$ is an abbreviation to a path $(a, \dots, b, \dots, c)$.

If nodes x, y are on some (directed or undirected) path p, we identify p_{xy} with the *open* subpath of p between nodes x and y, and by $[p_{xy}] = x - p_{xy} - y$ the *closed* subpath that includes the extreme vertices. Thus for example if $p = (a_1, a_2, a_3, a_4, a_5, a_6)$ we can write $p = a_1 - [p_{a_2a_4}] - a_5 - a_6$ or $p = a_1 - p_{a_1a_4} - a_4 - [p_{a_5a_6}]$. We denote by $|p|$ the number of edges in p, thus $|[p_{ab}]| = |p_{ab}|$.

Definition 1 ((Milchtaich, 2006; Holzman and Law-Yone, 2003)). *A* 2-terminal [directed] graph $G = \langle V, E, s, t \rangle$ *is a [directed] multigraph with no self-loops and two distinguished vertices* $s, t \in V$*, such that every vertex belongs to at least one [directed] simple* $s - t$ *path.*

A *forward-subtree* of a directed 2-terminal graph G is a subset of edges that form a directed tree with a single source. Similarly, A *backward-subtree* of G is a subset of edges that form a directed tree with a single target.

A directed 2-terminal graph that has no cycles is referred to as *TDAG* (2-Terminal Directed Acyclic Graph). We note that the vertices of a TDAG can always be sorted in increasing order (called their *topological order*) so that all edges (and thus all directed paths) are from v_i to v_j for some $j > i$. In particular s and t are the first and last vertices, respectively. Also note that in a TDAG all paths are simple. For technical reasons we also allow the degenerated TDAG with a single vertex (serving as both source and target).

2 Directed Graph Embeddings

Milchtaich (2006) gave a constructive definition for an undirected 2-terminal graph that is embedded in another graph.

Definition 2 ((Milchtaich, 2006)). *An undirected 2-terminal graph* H' *is* embedded *in* H *if* H' *is isomorphic to* H *or to a graph derived from* H *by applying the following operations any number of times in any order:*

Addition *The addition of a new edge joining two existing vertices;*

Terminal extension *The extension of a terminal vertex: addition of a new edge* e *joining* s *or* t *with another, new vertex, which becomes the new source or target, respectively.*

Subdivision *The subdivision of an edge* (a, b)*: its replacement by two edges* (a, x) *and* (x, b)*.*

A *subgraph* of H is a 2-terminal graph H' obtained only by removing edges and vertices. An easy observation is that if H' is a valid subgraph of H (i.e., H' is also a 2-terminal graph), then H' is embedded in H. Also, every embedding operation on H maintains the property that H is a 2-terminal graph. Thus the notion of embedding captures the idea that H' is "simpler" than H.

We would like to extend the definition of embedding to directed graphs. Unfortunately, a naïve extension of the definition will fail to meet our needs. First, if we allow arbitrary directed 2-terminal graphs then very basic properties of embedding are lost, e.g., adding an edge may result in an invalid 2-terminal graph (see Fig. 1a). Some of our definitions and results in this section only apply to 2-terminal graphs or to TDAGs (we state explicitly which ones). We thus modify the definition as follows. Note that it is not restricted to 2-terminal graphs.

Definition 3. *A directed graph* G' *is* d-embedded *in a directed graph* G *if* G' *is isomorphic to* G *or to a graph derived from* G *by applying the following operations any number of times in any order:*

Addition *The addition of a new edge* (a, b) *joining two existing vertices* $\{a, b\}$*, such that there is no path* $b - a$*;*

Forward split *A replacement of node* $a \neq t$ *by two nodes and an edge* (a, b)*, where* a *retains all incoming edges (and the name* a*), and* b *retains at least one outgoing edge;*

Backward split *A replacement of node* $b \neq s$ *by two nodes and an edge* (a, b)*, where* b *retains all outgoing edges, and* a *retains at least one incoming edge.*

Note that if $G' = \langle V', E' \rangle$ is d-embedded in $G = \langle V, E \rangle$, then we can identify for each node $x \in V'$ a corresponding node in V. This mapping may not be unique, as there may be several ways to obtain G from G'. However the mapping is unique for a given sequence of d-embedding operations.

(a) (b) Impossible subdivision (c) The Braess graph

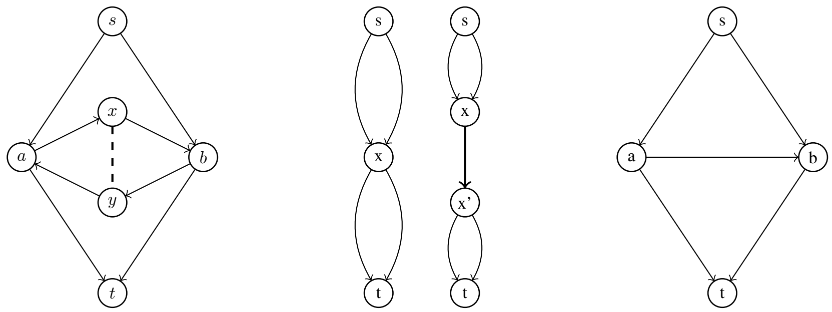

Figure 1: Examples. The graph in Fig. 1a is directed 2-terminal graph (solid edges only). Adding the dashed edge (x, y), regardless of its direction, results in an invalid 2-terminal graph. Fig 1b: The graph G' on the left is d-embedded in G on the right, as we can forward-split x into (x, x') (x retains all incoming edges, and x' retains at least one outgoing edge). However, there is no edge we can add or subdivide to get G from G'. The Braess graph G_B in on Fig. 1c.

Remark 1. *It is not hard to see that a subdivision of an edge (directed or undirected) can be replicated by splitting one of its end nodes, and a terminal extension can be replicated by splitting the terminal (backward split of s or forward split of t). We thus allow the operations of **edge subdivision** and **terminal extension** as d-embedding operations as well.*

We show that the class of TDAGs is closed under d-embeddings.

Lemma 1. *If G' is a TDAG and G' is d-embedded in G, then G is a TDAG.*

Proof. We need to show that after every single d-embedding operation on G', the new graph G is both a valid 2-terminal directed graph, and acyclic. For an addition step, since there is no path $b - a$, no cycle is formed. Clearly all existing edges and vertices are still part of some $s - t$ path. Consider some paths $p' = s - a$ and $p'' = b - t$. The new edge (a, b) is part of the path $p = s - p' - a - b - p'' - t$ in G. This path must be simple since G is a DAG.

For a (forward) split step, the path $p' = s - a - t$ becomes the path $p = [p'_{sa}] - b - [p'_{xt}]$, where x is the node following a in p', and p contains (a, b). The topological order of all vertices is maintained, where b is immediately after a. Similarly for backward split. □

If G' is a subgraph of G we say that G is a *supergraph* of G'. For a 2-terminal directed graph G, the graph G' is a *valid subgraph* of G if it is a subgraph of G and is isomorphic to a TDAG.

Lemma 2. *Let G be a TDAG. If G' is a valid subgraph of G then G' is d-embedded in G.*

Lemma 3. *If G, G' differ by a single forward split step of vertex a into (a, b), then there is a one to one mapping between paths in G' to paths in G.*

We get the next lemma as an immediate corollary, as addition steps never eliminate paths.

Lemma 4. *If G' is d-embedded in G, and G' contains a path from x to y, then G contains a path from x to y.*

Graph subdivision Holzman and Monderer (2015) consider directed 2-terminal networks, and define the notion of *graph subdivision*.

Definition 4 ((Holzman and Monderer, 2015)). *A directed graph G is a* subdivision *of a directed graph G' if G can be derived from G' by replacing every edge in G' with a path of one or more edge.*

Lemma 5. *Let G, G' be TDAGs. If G is isomorphic to a subdivision of a supergraph of G', then G' is d-embedded in G.*

Proof. We first recall that by Remark 1 subdivision of an edge is also a valid d-embedding step. Denote the supergraph of G' by $\hat{G}$. Thus G' is d-embedded in $\hat{G}$, and by Lemma 1, $\hat{G}$ is also a TDAG. Since $\hat{G}$ is a TDAG isomorphic to a subgraph of G, then by Lemma 2 $\hat{G}$ is d-embedded in G, and thus G' is d-embedded in G. □

The converse of Lemma 5 does not hold. That is, d-embedding allows for a richer set of operations than subdivision and adding edges. To see this, consider the graphs in Fig. 1b.

3 Graph minors

An undirected graph G' is a *minor* of graph G, if G' can be obtained from G by contracting edges, deleting edges, and deleting isolated vertices (Lovász, 2006). There are several extensions of this notion to directed graphs. One that is closest to our needs is the *butterfly minor* (Johnson et al., 2001).

Definition 5 ((Johnson et al., 2001))**.** *A graph G' is a* butterfly minor *of a directed graph G, if G' can be obtained from G by a sequence of the following local operations:*

Deletion *Deleting an edge (a, b);*

Backward contraction *Contracting an edge (a, b) where b has indegree 1;*

Forward contraction *Contracting an edge (a, b) where a has outdegree 1.*

The restricted contraction steps guarantee than no new paths or cycles are formed due to contraction, so the graph indeed becomes "simpler." However, The class of TDAGs is not closed under the butterfly minor operation (for example, it may leave an isolated node). We thus slightly modify it by restricting which edges may be deleted.

Definition 6. *A graph G' is a* directed minor *(or simply* a d-minor*) of a directed graph G, if G' can be obtained from G by a sequence of the following local operations:*

Deletion *Deleting an edge (a, b) where a has outdegree at least 2, and b has indegree at least 2.*

Backward contraction *Contracting an edge (a, b) where b has indegree 1;*

Forward contraction *Contracting an edge (a, b) where a has outdegree 1;*

Observation 6. *Let G be a directed graph. For any $e \in E$, there is a valid d-minor step on G that either deletes or contracts edge e.*

This is simply since every edge meets the premise of at least one d-minor operation. We next show that the class of TDAGs is closed under d-minor operations.

Lemma 7. *If G is TDAG and G' is a d-minor of G, then G' is a TDAG.*

Proof. Consider first a deletion step of edge (a, b). Clearly deletion does not add cycles, but we need to show that every edge $e \in E'$ is on some $s - t$ path in G'. Let p be some $s - t$ path in G containing e. If (a, b) is not in p then e is not affected. If (a, b) in p, consider the subpaths p_{sa} and p_{bt} of p. One of them contains e, w.l.o.g. p_{sa}. Note that since a has outdegree at least 2 in G, there is an edge (a, a') in E'. Moreover, this edge must belong to some path $p' = [p_{sa}] - a' - t$ in G, and this path may not contain (a, b) as this would mean that a appears twice in p'. The concatenation of p_{sa} and p'_{at} is a $s - t$ path that contains e in G'.

Next, consider a backward contraction step of edge (a, b). Every $s - t$ path p that goes through (a, b) in G, the path $s - p_{sa} - a - p_{at} - t$ exists in G', thus G' is a valid 2-terminal graph. We also note that there are no new paths: suppose that G' has a path $p' = x - a - y$ for some pair x, y. Then since before contraction b had indegree 1, the only incoming edge to b was from a, which means that path $[p'_{xa}] - b$ exists in G. Also, either the path $[p'_{by}]$ or $[p'_{ay}]$ is in G, and thus x, y are connected in G via the path $[p'_{xa}] - p'_{ay} - y$ or $[p'_{xa}] - [p'_{by}]$. $\square$

For directed 2-terminal graphs, the concepts of directed-minor and directed-embedding turn out to be equivalent.

Lemma 8. *Let G and G' be TDAGs. G' is d-embedded in G if and only if G' is a d-minor of G.*

Intuitively, addition and deletion operations cancel one another, as to split and contraction operations. This equivalence does not hold for general directed graphs, as added edges may not qualify for deletion (e.g. if we add an edge (a, b) where a has only incoming edges), and vice versa (if we remove an edge that is part of a cycle).

4 Parallel Width

Consider a directed 2-terminal graph $G = \langle V, E, s, t\rangle$. The following 3 definitions are due to (Meir and Parkes, 2017).

Definition 7. *A set of edges $S \subseteq E$ is* parallel *if there is some $S' \subseteq E$ s.t. $S \subseteq S'$, and S' is a minimal cut between s and t in the graph G.*

Definition 8. *A set of edges $S \subseteq E$ is* crossed *if there is a simple directed $s - t$ path p that contains S.*

Definition 9 (Parallel Width)**.** *The* parallel-width *of a directed 2-terminal graph, PW(G), is the size of the largest set $S \subseteq E$ that is both parallel and crossed.*

Intuitively, a parallel width of k means there are at least k source-target paths, and some additional path that edge-intersects all of them.

Example 9. *Consider the Braess graph in Fig. 1c. The minimal $s-t$ cuts in the graph are: $\{sa, sb\}$, $\{at, bt\}$, and $\{sb, ab, at\}$. Thus the set $\{sa, bt\}$ is both parallel and crossed, which means PW$(G_B) \geq 2$. The set $\{sa, at\}$ is crossed but not parallel; and $\{sa, sb, ab\}$ is neither. In fact, the* only *parallel set of size greater than 2 is $\{sb, ab, at\}$, which is not crossed, thus PW$(G_B) < 3$. We conclude that the parallel-width of the Braess graph is 2.*

For any 2-terminal graph G, we have $1 \leq \text{PW}(G) \leq |V| - 1$. The lower bound is since any single edge is both parallel and crossed, and the upper bound since there is no simple path of length $|V|$ or more.

Definition 10. *For any $k \geq 2$, we define the k-*crossed-parallel graph *$G_{CP(k)}$ as follows. $G = \langle V, E, s, t\rangle$, where $V = \{s, t, a_2, \ldots, a_k, b_1, \ldots, b_{k-1}\}$, and $E = \bigcup_{i=2}^{k-1}\{(s, a_i), (a_i, b_i), (b_i, t), (b_i, a_{i+1})\} \cup \{(s, b_1), (a_k, t)\}$.*

Definition 11. *A graph G is a* variant *of $G_{CP(k)}$ if we replace the edges $\{(s, a_i)\}_{i=2}^{k}$ with an arbitrary forward-subtree that respects the lexicographic order $(s, a_2, \ldots, a_k)$, and replace the edges $\{(b_i, t)\}_{i=1}^{k-1}$ with an arbitrary backward-subtree that respects the lexicographic order $b_1, \ldots, b_{k-1}, t$ (note the set of vertices remains the same).*

See Figure 2 for examples.

The 2-crossed-parallel graph is the Braess graph G_B, which has only one variant. There are other generalizations of the Braess graph but we are unaware of one that coincides with ours.[1] We should note that an example of a $G_{CP(k)}$ graph appears in Nikolova and Stier-Moses (2015) (Figure 4 there) without a formal definition

[1]Lianeas et al. (2015), for example, generalize the Braess graph to a family of graphs with long alternating paths. All these graphs have parallel-width 2, whereas $G_{CP(k)}$ graphs have alternating paths of length at most 2.

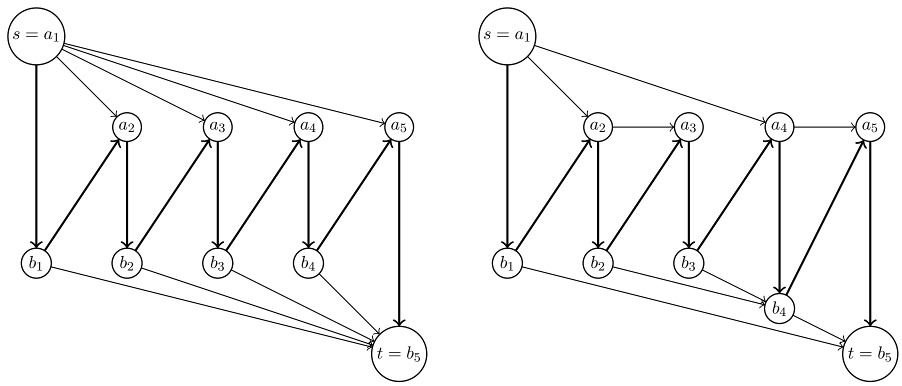

Figure 2: The left figure is the graph $G_{CP(5)}$, and the right figure is a variant of it. For convenience, the long path in each graph appears in thick lines, and the forward- and backward-trees in thin lines.

or analysis, and is also used in (**?**) to derive an example where few malicious players can hurt the others. The parallel-width of the k-crossed-parallel graph is exactly k, where $\{(s, b_1), (a_2, b_2), \ldots, (a_{k-1}, b_{k-1}), (a_k, t)\}$ are the parallel edges.

Lemma 10. *If $PW(G') \geq k$ and G' is d-embedded in G, then $PW(G) \geq k$.*

Proof. Denote graph $G' = \langle V', E', s', t' \rangle$. Denote by S' some set of k edges in G' that are parallel and crossed, and by C' some minimal cut containing S'. It is sufficient to show that $PW(G) \geq k$ for G obtained after a single embedding operation (add edge or split node).

- add edge e: Either C' is still a minimal cut, or $C' \cup \{e\}$ is. Also, the path containing S' still exists after adding an edge.
- split node: By Lemma 3, G and G' have exactly the same set of paths, and thus C' is still a minimal cut in G. For the same reason, there is a path containing S' in G.

In either case we get that the set S' is both parallel and crossed in the new graph, so $PW(G) \geq k$. □

4.1 Characterization of graphs with low parallel width

Theorem 11. *Let G be a TDAG, and let $k \geq 2$. The following conditions coincide.*

1. *PW$(G) \geq k$;*
2. *Some variant of $G_{CP(k)}$ is a d-minor of G;*
3. *Some variant of $G_{CP(k)}$ is d-embedded in G.*

Proof. "1 ⇒ 2": Consider the graph G. Suppose that PW$(G) \geq k$, then there is a set $S = \{e_1, \ldots, e_k\}$ that is part of a minimal cut C between s and t. For every $i \leq k$, there must be some path p_i from s to t such that $e_i \in p_i$, and $e_j \notin p_i$ for all $j \neq i$, otherwise we could drop e_i from S and still get a $s - t$ cut $C \setminus \{e_i\}$. Denote $e_i = (a_i, b_i)$. If there are several such paths, then p_i is the path among them that has maximum edge overlap with the path p_{i-1}.

Also, by definition of the parallel width there is a *simple* $s - t$ path p' containing S, w.l.o.g. in that lexicographic order. If there are several such paths, p' is the one maximizing edge overlap with the union of $\{p_i\}_{i=1}^k$.

We now describe a series of d-minor steps on G that will result in a variant of $G_{CP(k)}$. Delete all edges and vertices that are not part of the paths p' or p_i for some $i \leq k$. This leaves us with a graph G' that is a subgraph of G and has the following form: The union of $\{p_i\}$ can be written as the union of S, a forward-subtree T_s whose source is s and whose leafs are $\{a_i\}_{i=1}^k$ (unless $a_1 = s$); and a backward-subtree T_t whose target is t and whose sources are $\{b_i\}_{i=1}^k$ (unless $b_k = t$). Since G' is a valid subgraph of G, by Lemmas 2 and 8 it is also a d-minor of G.

p' is composed of a sequence of subpaths between vertices $s, y_1, x_2, y_2 \ldots, x_{k-1}, y_{k-1}, x_k, t$, where each x_i is the convergence point of p_i and p' on tree T_s, such that x_i is an ancestor (or coincides with) of a_i Similarly, $\{y_i\}_{i=1}^{k-1}$ are on the backward-subtree T_t, where y_i is the divergence point of p' and p_i. Denote by $A_i \subseteq \{a_2, \ldots, a_k\}$ all leafs of the subtree of T_s rooted by x_i, and by $B_i \subseteq \{b_1, \ldots, b_{k-1}\}$ all roots of the subtree of T_t whose leaf is y_i. In particular, $a_i \in A_i$, and $a_j \notin A_i$ for $j < i$, as otherwise there is a cycle $x_i - a_j - b_j - y_j - x_i$. Likewise, $b_i \in B_i$ and $b_j \notin B_i$ for $j > i$.

Note that the indegree of all nodes in T_s is 1, except for $\{x_i\}_{i=2}^k$ whose indegree is 2 (one edge from the parent in T_s, and one from the predecessor node on p'), and s whose indegree is 0. We thus backward-contract all edges in T_s that do not point to some x_i. This leaves us with a forward-subtree $\hat{T}_s$:

- The root of $\hat{T}_s$ is $s = x_1$, and its nodes are $\{x_i\}_{i=2}^k$;
- Each path $x_i - a_i$ is contracted to a single node $x_i = a_i$;
- The subtree rooted by x_i in T_s is contracted to a path in $\hat{T}_s$ containing all nodes of A_i in their lexicographic order.

Similarly, $\{y_i\}$ are the only nodes in T_t whose outdegree is > 1. After forward-contracting all edges of T_t not originating in some y_i, we get a backward-subtree $\hat{T}_t$ over nodes $\{y_i\}_{i=1}^{k-1}$ whose leaf is $t = y_k$. Each subtree whose leaf is y_i is contracted to a tree over nodes B_i maintaining their order (i.e., children have higher index than their parent). Similarly, $\{y_i\}$ are the only nodes in T_t whose outdegree is > 1. After forward-contracting all edges of T_t not originating in some y_i, we get a backward-subtree $\hat{T}_t$ over nodes $\{y_i\}_{i=1}^{k-1}$ whose leaf is $t = y_k$. Each subtree whose leaf is y_i is contracted to a tree over nodes B_i maintaining their order (i.e., children have higher index than their parent).

The last step is to contract every subpath $[p'_{y_i, x_{i+1}}]$ to a single edge $(y_i, x_{i+1} = a_{i+1})$. Denote the union of these edges by $\hat{X}$, so that $S \cup \hat{X}$ is the path we got after contracting p'.

We get that the contracted graph $\hat{G}' = S \cup \hat{X} \cup \hat{T}_s \cup \hat{T}_t$ is isomorphic to a variant of $G_{CP(k)}$. More specifically, s and t are isomorphic to themselves, each x_i for $i = 2, \ldots, k$ in $\hat{G}$ is isomorphic to a_i in $G_{CP(k)}$, and each y_i for $i = 1, \ldots, k-1$ in $\hat{G}'$ is isomorphic to b_i in $G_{CP(k)}$. For each $i = 2, \ldots, k$, let j be the maximal index such that x_j is an ancestor of x_i in T_s. If such j exists, then the parent of a_i in $\hat{G}'$ is a_j, and otherwise its parent is $s = a_1$. The parent of a_i in $G_{CP(k)}$ is the closest ancestor x_j of the node x_i in T_s (and similarly for the child of b_i).

"$2 \Rightarrow 3$": Follows directly from Lemma 8.

"$3 \Rightarrow 1$": Follows directly from Lemma 10, since $PW(G_{CP(k)}) = k$, and likewise for any variant of $G_{CP(k)}$. □

Since $G_{CP(k)}$ has $2k$ vertices, we get the following bound:

Corollary 12. *For any TDAG $G = \langle V, E \rangle$, $PW(G) \leq \frac{|V|}{2}$.*

4.2 Series-parallel graphs

Series-parallel 2-terminal graphs have been long studied in contexts such as electric circuits (Duffin, 1965), complexity of graph algorithms (Takamizawa et al., 1982), and also routing games (Milchtaich, 2006; Fotakis and Spirakis, 2007; Epstein et al., 2009). Whereas most papers assume undirected graphs, there is a very similar definition for directed graphs due to (Jakoby et al., 2006), thus we bring them together.

Definition 12 ((Eppstein, 1992; Jakoby et al., 2006))**.** *A* [directed] series-parallel graph *is a 2-terminal graph $\langle V, E, s, t\rangle$, and is either a single edge (s,t), or is composed recursively by one of the two following steps:*

Serial composition *combine two [directed] 2-terminal graphs $\langle V_1, E_1, s_1, t_1\rangle, \langle V_2, E_2, s_2, t_2\rangle$ serially by merging t_1 with s_2;*

Parallel composition *combine two [directed] 2-terminal graphs $\langle V_1, E_1, s_1, t_1\rangle, \langle V_2, E_2, s_2, t_2\rangle$ in parallel by merging s_1 with s_2, and t_1 with t_2.*

Our second main result in this section is showing that directed series-parallel graphs (DSP) characterize exactly the 2-terminal graphs with parallel-width of 1 (Theorem 19 below). We prove this using a known characterizations of directed and undirected series-parallel graphs, thereby also shed some light on the connections between embeddings, d-embeddings, and subdivisions.

Directed graph subdivisions

Proposition 13 ((Holzman and Monderer, 2015))**.** *Let G be a 2-terminal directed graph. Then G is a DSP if and only if G has no subgraph that is isomorphic to a subdivision of the Braess graph G_B.*

Undirected graph embeddings

Proposition 14 ((Duffin, 1965; Milchtaich, 2006))**.** *For an undirected 2-terminal graph H, the following conditions coincide: 1. H is series-parallel; 2. for every pair of distinct vertices u and v, if u precedes v in some path p containing both vertices, then u precedes v in all such paths; and 3. the (undirected) Braess graph[2] $\overline{G}_B$ is not embedded in H.*

We can rephrase condition (2) above, to get two more equivalent conditions that will be useful.

Lemma 15. *For an undirected 2-terminal graph H, the following conditions coincide: 1. H is series-parallel; 4. there are no two paths in H that go through the same edge in opposite directions; and 5. there is a unique 2-terminal directed graph G such that $\overline{G} = H$.*

Proof. We use the equivalence of conditions (1) and (2) from Proposition 14. That condition (2) entails (4) is obvious, as the existence of such a bidirectional edge (a, b) entails the existence of a pair $\{u = a, v = b\}$. (4) entails (5): Suppose there are two directed 2-terminal graphs G, G' such that $\overline{G} = \overline{G}' = H$. This means there is some edge $(a, b) \in E$ such that $(b, a) \in E'$. Since the graphs are 2-terminal, there is a directed $s - t$ path p in G that contains (a, b), and likewise a directed $s - t$ path q in G' that contains (b, a). Both of p, q are valid $s - t$ paths in H that go through (a, b) in opposite directions.

Finally, (5) entails (2), since if there is a unique way to direct the edges of H, then every $s - t$ path in H uniquely determines an ordering of all vertices along the path. □

Unfortunately, since d-embedding allows for a different set of operations than "standard" undirected embedding (as per Def. 2), we cannot use Prop. 14 directly. However we can use an alternative definition of embeddings due to (Epstein et al., 2009). We rephrase their definition to make it more similar in style to our other definitions.[3]

Definition 13 ((Epstein et al., 2009), rephrased)**.** *A 2-terminal undirected graph H' is* split-embedded *(or s-embedded) in H if H' is isomorphic to H or to a graph derived from H by applying the following operations any number of times in any order:*

Addition *The addition of a new edge joining two existing vertices;*

[2] This graph is also known as the *Wheatstone network* (Milchtaich, 2006).

[3] (Epstein et al., 2009) used the reversed operations of edge removal and edge contraction, as in undirected minors.

(a) G' and G (b) $\overline{G}'$ and $\overline{G}$

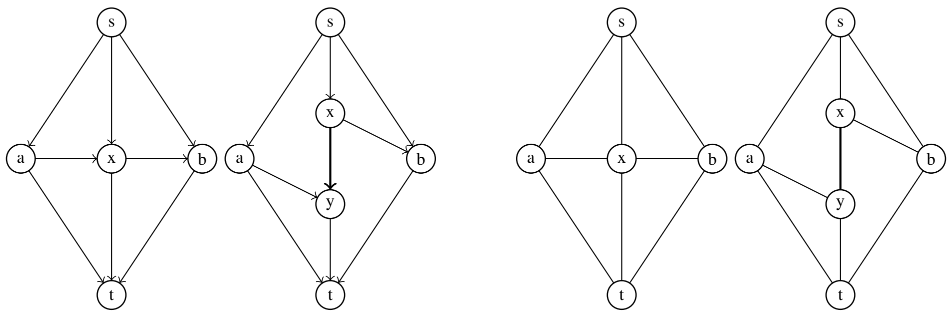

Figure 3: The graph G' is not d-embedded in G, since x retains an outgoing edge and y retains an incoming edge (indeed, G' has a path $a - b$ that does no exist in G). In contrast, splitting x to the edge (x, y) is a valid s-embedding step in $\overline{G}'$.

Terminal extension *The extension of a terminal vertex: addition of a new edge e joining s or t with another, new vertex, which becomes the new source or target, respectively.*

Split *A replacement of a non-terminal node a by two nodes a' and b' and an edge (a', b'), such that each new node retains at least one edge of a.*

Interestingly, the characterization of series-parallel networks in (Milchtaich, 2006) still holds with this more lax definition of s-embedding, without changing Milchtaich's proof. From that proof and Lemma 15 we get the following. For completeness, the full proof is in the appendix.

Proposition 16. *For an undirected 2-terminal graph H, the following conditions coincide:*

1. *H is series-parallel;*
2. *for every pair of distinct vertices u and v, if u precedes v in some path p containing both vertices, then u precedes v in all such paths;*
3. *the (undirected) Braess graph $\overline{G}_B$ is not **s-embedded** in H;*
4. *there are no two paths in H that go through the same edge in opposite directions;*
5. *there is a unique 2-terminal directed graph G such that $\overline{G} = H$.*

There is a natural correspondence between d-embeddings and s-embeddings. This is since we can replicate the sequence of d-embedding operations on G' with s-embedding operations on $\overline{G}'$.

Observation 17. *If G' is d-embedded in G then $\overline{G}'$ is s-embedded in $\overline{G}$.*

This is obvious since for every d-embedding step on G' we can apply a corresponding s-embedding step on $\overline{G}'$. However, the other direction does not hold in general. Fig. 3 shows a counter example.

Lemma 18. *G is a DSP if and only if $\overline{G}$ is series-parallel.*

Proof. Assume by induction that this is true for any graph with k edges (the base case of $k = 1$ is obvious). If G is a DSP then it is constructed of two smaller DSPs G_1, G_2 in series [in parallel]. By induction, $\overline{G}_1, \overline{G}_2$ are series-parallel. We can construct $\overline{G}$ by joining $\overline{G}_1, \overline{G}_2$ in series [resp., in parallel].

In the other direction, suppose that $H = \overline{G}$ is series-parallel. Consider the last (series/parallel) construction step of H from H_1, H_2. Direct the edges of H_1, H_2 to obtain directed graphs G_1, G_2. By induction, G_1, G_2 are DSPs, and thus joining them (in series/parallel) results in a DSP G_{1+2} such that $\overline{G}_{1+2} = H = \overline{G}$. Finally, by conditions (1) and (5) of Lemma 15 there is only one way to direct the edges of H from s to t and thus $G = G_{1+2}$. □

We can now turn to prove a characterization of DSPs in terms of their parallel-width.[4] It will follow as an easy corollary of the previous results in the paper.

Theorem 19. *Let G be a TDAG, and let $k \geq 2$. The following conditions coincide:*

1. *G is a directed series-parallel graph;*
2. *The directed Braess graph G_B is not d-embedded in G;*
3. *$PW(G) = 1$;*

Proof. "1 $\Rightarrow$ 2": Suppose that G_B is d-embedded in G. Then by Observation 17, $\overline{G}_B$ is s-embedded in $\overline{G}$. Due to Prop. 16 (conditions (1) and (3)), this means that $\overline{G}$ is not series-parallel. However this means that G is not a DSP by Lemma 18.

"2 $\Rightarrow$ 1": Suppose that G is not a DSP. Then by Prop. 13, G is isomorphic to a subdivision of a supergraph of the Braess graph G_B. Then, by Lemma 5, G_B is d-embedded in G.

"2 $\iff$ 3": Follows directly from Theorem 11.

□

5 Discussion

There are by now many different variations of operations that can be used to obtain "simple" graphs that capture the essential forbidden properties of large classes of graphs: minors, embeddings, subdivisions, etc. These operations should be rich enough to allow for a small set of forbidden graphs, but restricted enough to only capture the intended class.

For example, s-embedding, which is richer than other embeddings for undirected graphs, allows for a characterization of extension-parallel networks with just 1 forbidden graph instead of 3 (see (Milchtaich, 2006; Epstein et al., 2009)). One can check that applying s-embeddings to other characterizations can reduce the number of required forbidden graphs from 5 (in Lemma 2 of (Milchtaich, 2015)) to 2, or from 9 (in Proposition 2 of (Acemoglu et al., 2016)) to 2. Also, for each of the graphs K_5 and $K_{3,3}$ used in the characterization of planar graphs (Wagner, 1937), 9 more graphs would be required if we only allow edge subdivision rather than node splits (the reverse operation of edge contraction).

We hope that the notions of d-embedding and d-minor (which coincide for directed 2-terminal acyclic graphs) will turn out as useful ones, beyond the applications demonstrated in the paper. Some more indication to their potential is that they (somewhat) simplify known characterizations. For example, multiextension-parallel graphs characterizations with subdivision graphs require 2 forbidden graphs (Holzman and Monderer, 2015), but one of them is d-embedded in another other, and it is easy to verify that this 2 graphs suffice to characterize the class with d-embeddings. To characterize acyclic graphs, only one graph is required (a bidirectional edge). It is interesting whether d-embeddings or d-minors can be used to characterize other classes of directed graphs, such as graphs with bounded triangular width (Meer, 2011), D-width (Safari, 2005), directed path-width (Barát, 2006), or DAG-width (Obdržálek, 2006).

Finally, given there is the question of whether a directed graph version of Wagner's conjecture (Robertson and Seymour, 2004) holds for d-minors or d-embeddings.

Acknowledgments

The author thanks Ron Holzman, Pieter Kleer, Igal Milchtaich, Nicolás Stier-Moses, Evdokia Nikolova, David Parkes and Éva Tardos for helpful discussions and comments.

[4] In (Holzman and Law-Yone, 2003) there is a similar characterization of *extension-parallel graphs*, as graphs where there are no two edges e, e' and three paths such that one uses only e, one uses only e' and the third uses both. Note that this condition implies a parallel-width of 1 but not vice versa, as such e, e' may not be contained in any minimal cut.

Bibliography

Daron Acemoglu, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar. 2016. Informational Braess' Paradox: The Effect of Information on Traffic Congestion. *arXiv preprint arXiv:1601.02039* (2016).

Moshe Babaioff, Robert Kleinberg, and Christos H Papadimitriou. 2007. Congestion games with malicious players. In *Proceedings of the 8th ACM conference on Electronic commerce*. ACM, 103–112.

János Barát. 2006. Directed path-width and monotonicity in digraph searching. *Graphs and Combinatorics* 22, 2 (2006), 161–172.

R.J Duffin. 1965. Topology of series-parallel networks. *J. Math. Anal. Appl.* 10, 2 (1965), 303 – 318.

David Eppstein. 1992. Parallel recognition of series-parallel graphs. *Information and Computation* 98, 1 (1992), 41–55.

Amir Epstein, Michal Feldman, and Yishay Mansour. 2009. Efficient graph topologies in network routing games. *Games and Economic Behavior* 66, 1 (2009), 115–125.

Dimitris Fotakis and Paul Spirakis. 2007. Cost-balancing tolls for atomic network congestion games. *Internet and Network Economics* (2007), 179–190.

Ron Holzman and Nissan Law-Yone. 2003. Network structure and strong equilibrium in route selection games. *Mathematical Social Sciences* 46, 2 (2003), 193–205.

Ron Holzman and Dov Monderer. 2015. Strong equilibrium in network congestion games: increasing versus decreasing costs. *International Journal of Game Theory* 44, 3 (2015), 647–666.

Andreas Jakoby, Maciej Liśkiewicz, and Rüdiger Reischuk. 2006. Space efficient algorithms for directed series–parallel graphs. *Journal of Algorithms* 60, 2 (2006), 85–114.

Thor Johnson, Neil Robertson, Paul Seymour, and Robin Thomas. 2015. Excluding A Grid Minor In Planar Digraphs. *arXiv preprint arXiv:1510.00473* (2015).

Thor Johnson, Neil Robertson, Paul D Seymour, and Robin Thomas. 2001. Directed tree-width. *Journal of Combinatorial Theory, Series B* 82, 1 (2001), 138–154.

Ken-ichi Kawarabayashi and Bojan Mohar. 2007. Some recent progress and applications in graph minor theory. *Graphs and combinatorics* 23, 1 (2007), 1–46.

Shiva Kintali and Qiuyi Zhang. 2013. Forbidden directed minors and Kelly-width. *arXiv preprint arXiv:1308.5170* (2013).

Casimir Kuratowski. 1930. Sur le probleme des courbes gauches en topologie. *Fundamenta mathematicae* 15, 1 (1930), 271–283.

Thanasis Lianeas, Evdokia Nikolova, and Nicolas E Stier-Moses. 2015. Asymptotically tight bounds for inefficiency in risk-averse selfish routing. *arXiv preprint arXiv:1510.02067* (2015).

László Lovász. 2006. Graph minor theory. *Bull. Amer. Math. Soc.* 43, 1 (2006), 75–86.

Klaus Meer. 2011. An extended tree-width notion for directed graphs related to the computation of permanents. *Computer Science–Theory and Applications* (2011), 247–260.

Reshef Meir and David C. Parkes. 2017. Playing the Wrong Game: Smoothness Bounds for Congestion Games with Heterogeneous Biases. *CoRR* abs/1411.1751v5 (2017).

Igal Milchtaich. 2006. Network topology and the efficiency of equilibrium. *Games and Economic Behavior* 57 (2006), 321–346.

Igal Milchtaich. 2015. Network topology and equilibrium existence in weighted network congestion games. *International Journal of Game Theory* 44, 3 (2015), 515–541.

Evdokia Nikolova and Nicolas E Stier-Moses. 2015. The burden of risk aversion in mean-risk selfish routing. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM, 489–506.

Jan Obdržálek. 2006. DAG-width: connectivity measure for directed graphs. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*. Society for Industrial and Applied Mathematics, 814–821.

Neil Robertson and Paul D. Seymour. 1986. Graph minors. II. Algorithmic aspects of tree-width. *Journal of algorithms* 7, 3 (1986), 309–322.

Neil Robertson and Paul D Seymour. 2004. Graph minors. XX. Wagner's conjecture. *Journal of Combinatorial Theory, Series B* 92, 2 (2004), 325–357.

Mohammad Ali Safari. 2005. D-width: A more natural measure for directed tree width. In *MFCS*, Vol. 5. Springer, 745–756.

Kazuhiko Takamizawa, Takao Nishizeki, and Nobuji Saito. 1982. Linear-time computability of combinatorial problems on series-parallel graphs. *Journal of the ACM (JACM)* 29, 3 (1982), 623–641.

Klaus Wagner. 1937. Über eine Eigenschaft der ebenen Komplexe. *Math. Ann.* 114, 1 (1937), 570–590.

A Omitted proofs

Lemma 2. *Let G be a TDAG. If G' is a valid subgraph of G then G' is d-embedded in G.*

Proof. We prove by induction on the difference $|E| - |E'|$. Let $e = (a, b)$ be an edge in $E \setminus E'$. As e is part of path $p = s - a - b - t$ in G, let x be the last node on the subpath $[p_{sa}]$ that is also in V' (such a node must exist since $s \in V'$). Also let y be the first node on $[p_{bt}]$ that is in V'. Let k be the length of the path $[p_{xy}]$. We perform the following d-embedding steps: add an edge (x, y) to G', and subdivide it $k - 1$ times. The result G'' is a valid subgraph of G, that has strictly more than $|E'|$ edges. Thus by induction there is a sequence of d-embedding steps from G'' to G, and hence from G' to G. □

Lemma 3. *If G, G' differ by a single forward split step of vertex a into (a, b), then there is a one to one mapping between paths in G' to paths in G. Specifically for any path $p' = x - y$ in G' there is a path p in G as follows:*

- *If $a \notin p'$ then $p = p'$;*
- *If $p' = x - a - y$ then either $p = p'$ or $p = [p'_{xa}] - [p'_{by}]$.*

For any path $p = x - y$ (where $x \neq b$) in G, there is a path p' in G' as follows:

- *If $b \notin p$ then $p' = p$;*
- *If $p = x - a - b - y$ then $p' = x - p_{xa} - a - p_{ay} - y$;*
- *It is not possible that p contains b and not a.*

(a similar mapping holds after a backward-split).

Proof. Consider a forward split of node a in G' to edge (a, b). If the $x - y$ path does not go through a, then it still exists in G. Otherwise there is a path $p' = x - a - y$ in G'. Let z be the vertex immediately following a on this path (possibly $y = z$). After the split the graph G contains either the edge (a, z), in which case it contains the simple path $[pxa] - [p_{zy}] = p'$; or G contains the edge (b, z), and thus the simple path $[p'_{xa}] - b - [p'_{zy}] = [p'_{xa}] - [p'_{by}]$.

In the other direction, we only need to show that there is no path in G that contains only b. This follows immediately from the fact that a retains all incoming edges, so any path through b must also go through a. □

Lemma 8. *Let G and G' be TDAGs. G' is d-embedded in G if and only if G' is a d-minor of G.*

Proof. By induction, it is sufficient to show this for G', G that differ by a single d-embedding or d-minor operation. "⇒" There are 3 cases, depending on the embedding operation:

1. The addition of edge (a, b) to G' can be reversed by deleting the same edge from G. Note that $b \neq s$ as otherwise there is a path in G' from $b = s$ to a, and similarly $a \neq t$. Thus a has outdegree at least 1 in G' and at least 2 in G. Similarly, b has indegree at least 2 in G, and thus deleting the edge (a, b) is a valid d-minor step.

2. Suppose that a vertex a in G' is split to $\{a, b\}$ with a forward split. Then since a retains all incoming edges, b has a single incoming edge (a, b) in G. Thus we can contract the edge (a, b) in G using backward contraction.

3. Similarly, a backward split can be reversed with a forward contraction.

"⇐" There are 3 cases, depending on the d-minor operation:

1. If the edge (a,b) is deleted from G, then since G is acyclic there is no path $b-a$. Thus adding (a,b) to G' is a valid d-embedding step.
2. Suppose that the edge (a,b) in G is backward-contracted to some vertex x in G'. This means that b has a single incoming edge. Thus all edges incoming to the pair $\{a,b\}$ are leading to a. Let $R(a), R(b)$ be the out-neighbors of a and b in G, respectively. Then by forward-splitting node x in G' and split the outgoing edges of x according to $R(a)$ and $R(b)$, we get the graph G^i.
3. Similarly, forward contraction can be reversed with backward split.

□

Proposition 16. *For an undirected 2-terminal graph H, the following conditions coincide:*

1. *H is series-parallel;*
2. *for every pair of distinct vertices u and v, if u precedes v in some path p containing both vertices, then u precedes v in all such paths;*
3. *the (undirected) Braess graph $\overline{G}_B$ is not **s-embedded** in H;*
4. *there are no two paths in H that go through the same edge in opposite directions;*
5. *there is a unique 2-terminal directed graph G such that $\overline{G} = H$.*

Proof. By Prop. 14 and Lemma 15, and since any embedding operation is also an s-embedding operation, it only remains to show that condition (3) is entailed by one of the others. The undirected Braess graph $\overline{G}_B$ has an edge (a,b) and paths p,q that go through (a,b) in opposite directions. We show that this property remains after any s-embedding operation, which means that condition (4) entails (3).

Note that $\{a,b\} \cap \{s,t\} = \emptyset$. Denote by x_a, x_b the nodes that precedes a / succeeds b in p, respectively. Denote by y_a, y_b the nodes that succeeds a / precedes b in q. For addition and terminal extension this is obvious, as well as for any split operation of a node in $V \setminus \{a,b\}$. Suppose that a is split into (a', a''). The new graph contain either the edge (a', b) or (a'', b). Assume w.l.o.g. the former, and we will show that there are still $s-t$ paths through (a', b) in both directions. If a' retains the edge to x_a, then there is a path $p' = [p_{sx_a}] - a' - [p_{bt}]$, and otherwise, a'' retains the edge to x_a and there is a path $p' = [p_{sx_a}] - a'' - a' - [p_{bt}]$. The construction of a path q' through (b, a') is similar. □

B Additional proof of Theorem 19

Lemma 20. *Let $G = \langle V, E, s, t\rangle$ be a directed 2-terminal graph. A subset $S \subseteq E$ is contained in a minimal $s-t$ cut in G if and only if there is a valid subgraph $G' = \langle V', E', s, t\rangle$ of G such that S is a minimal $s-t$ cut in G'.*

Proof. "⇒" Suppose S is contained in a minimal cut C of G. Then for each $e \in S$ there is an $s-t$ path p_e that contains e but not any other edge from C. Let G' be the union of these paths $\bigcup_{e\in S} p_e$ (note that this is a valid subgraph of G), thus $E' \subseteq E \setminus (C \setminus S)$. Any $s-t$ path q in G' is also an $s-t$ path in G, and thus must contain some edge $e_q \in C$. But $C \cap E' = S$ and thus $e_q \in S$, which means that S is an $s-t$ cut in G'. For any $S' \subsetneq S$, we can take some $e' \in S \setminus S'$ and then the path $p_{e'}$ is an $s-t$ path in G' that does not contain any edge in S', meaning that S is a minimal cut in G'.

"⇐" Suppose S is a minimal cut in some subgraph G'. If S is a cut in G then it must be minimal and we are done. Otherwise, let $C' = S \cup (E \setminus E')$, and note that C' is an $s-t$ cut (not necessarily minimal) in G.

As long as C' is not a minimal cut, there is at least one edge e' we can remove such that $c' \setminus \{e'\}$ is still a cut in G. We thus remove such edges from C' in arbitrary order until we get a minimal cut C in G. We argue that $S \subseteq C$, since for each $e \in S$ there is a, $s-t$ path p_e in G' that does not contain any edge in $S \setminus \{e\}$. Further, since $p_e \subseteq E'$, p_e does not contain any edge in $C \setminus E' = C \setminus S$, then $p_e \cap C = \{e\}$, meaning e must be part of the cut C. □

Theorem 19. *Let G be a TDAG, and let $k \geq 2$. The following conditions coincide:*

1. *G is a directed series-parallel graph;*
2. *The directed Braess graph G_B is not d-embedded in G;*
3. *PW$(G) = 1$;*

The difficult part is to show "$3 \Rightarrow 1$", and the proof idea is to assume that the entailment holds for small graphs by induction. Then, for every larger graph, we show that it either contains the Braess graph as a minor (and thus have $PW(G) \geq 2$), or that there is an edge we can delete/contract, to get a smaller graph G' that also violates the entailment. This would contradict the induction hypothesis, so all graphs must respect the entailment. The proof itslefs requires a lengthy case analysis of the contracted edge.

Proof of "$3 \Rightarrow 1$". Assume that $G = \langle V, E, s, t\rangle$ is the smallest graph violating the entailment. That is, the graph with the minimal number of edges such that $PW(G) = 1$ but G is not a DSP. Then every subgraph of G that is a valid TDAG must be a DSP (this will be referred to as our *minimality assumption*).

By Lemma 18, $\overline{G}$ is not series parallel. Thus by conditions (1) and (4) of Prop. 16 there is an edge $e \in E$ that belongs to two undirected $s-t$ paths in $\overline{G}$ that pass through e in opposite directions. Since G is a TDAG, there is at least one directed $s-t$ path p that goes through e and we denote by q an *undirected* $s-t$ path that goes through e in the opposite direction. We select q such that $|q \setminus p|$ is minimal. Denote by a the last node in $p \cap q$ that weakly precedes e on p, and by b the first node in $p \cap q$ that weakly succeeds e. Thus $e \in [p_{ab}]$, and q_{ba} is p_{ab} in the opposite direction. More generally, p can be written as the directed path $p = s - p_{sa} - a - p_{ab} - b - p_{bt} - t = [p_{sa}] - p_{ab} - [p_{bt}]$, and q can be written as the *undirected path* $q = s - q_{sb} - b - q_{ba} - a - q_{at} - t = [q_{sb}] - q_{ba} - [q_{at}]$.

We now prove the following claim: $[p_{ab}] = [q_{ba}]$ is the *only* intersection of p and q.

To prove the claim, we first show that $q_{at} \cap p_{sa} = \emptyset$. Otherwise, let a^* be the node preceding a on p. By our selection of a, we know that $a^* \notin q$. Thus there is some $x \in p_{sa^*} \cap q_{at}$ (select the closest x to a under this condition), and $[q_{at}] = a - q_{ax} - x - q_{xt}$. But then there is a path $q' = [q_{sa}] - p_{ax} - [q_{xt}]$ that contains e and $|q' \setminus p| = |q \setminus p| - |q_{ax}| < |q \setminus p|$, in contradiction to our selection of q. Similarly, $q_{ab} \cap p_{bt} = \emptyset$.

Next, by the way we select q, the only possible intersection of p, q except the subpath $[p_{ab}] = [q_{ba}]$, is if p and q share a prefix (i.e., a subpath $[p_{ss'}] = [p_{sa}] \cap [q_{sb}]$) and/or a suffix (a subpath $[p_{t't}] = [p_{bt}] \cap [q_{at}]$). See Fig. 4(4a) for an example.

We argue that sharing a prefix or a suffix leads to a contradiction of our minimality assumption. Indeed, suppose that there is a nonempty shared prefix $[p_{ss'}]$. Select an arbitrary edge $e' \in p_{ss'}$. By Observation 6 we can either delete or contract the edge e' and obtain a d-minor G'' of G. If e' is contracted, then G'' still contains paths p' and q' (just p and q with one contracted edge) that pass through p_{ab} in opposite directions. Thus again by conditions (1) and (4) of Prop. 16, $\overline{G}''$ is not series-parallel, which by Lemma 18 means that G'' is not a DSP, i.e., a contradiction.

If e' is deleted, we need to work a bit harder, but we can still find such paths p' and q'. Note that there must be some directed path h in G'' from s to s', which due to acyclicity may not intersect path $p_{s't}$. Let s'' be the first node in h that is also in $[q_{s'b}]$ (must exist and cannot be b). Then $p' = [h_{ss''}] - h_{s''s'} - [p_{s't}]$ and $q' = [h_{ss''}] - q_{s''b} - [q_{bt}]$. The proof that p, q have no shared suffix is symmetric. This completes the proof of the claim.

(a) A graph G with shared prefix

(b) The graph G'

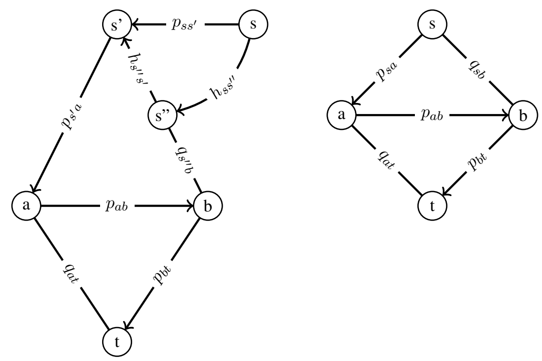

Figure 4: An arrow in the figure represents a directed path in the graph G, where as a line without arrow indicates an undirected path in $\overline{H}$. On the left we see an example of a graph G (or part of a graph), where paths p and q share a prefix. On the right we see the structure of $G' = \langle V, p \cup q, s, t\rangle$, after proving that $p \cap q = [p_{ab}]$.

Denote $G' = \langle V, p \cup q, s, t\rangle$. Note that $\overline{G}'$ is a valid 2-terminal undirected subgraph of $\overline{G}$, but G' itself may not be a valid 2-terminal directed graph (as some edges in q_{sb} or q_{at} may be reversed). The structure of G' is shown in Fig. 4(4b).

Note first that if $p \cup q = \overline{E}$ then q_{sb} and q_{at} are directed paths in G. In particular, take any edges $e_1 \in p_{sa}$ and $e_2 \in p_{bt}$, then $S = \{e_1, e_2\}$ is both crossed (due to path p) and parallel ($S' = \{e_1, e_2, (a, b)\}$ is a minimal cut in G' and thus by Lemma 20 contained in a minimal cut of G). Another way to see it is that $\langle V, p \cup q, s, t\rangle$ is a valid subgraph of (and thus d-embeded in) G, and that the Braess graph G_B is d-embedded in G' (and thus in G). This would mean that $PW(G) \geq 2$ in contradiction to the premise.

Thus there is some edge $(x, y) \in E \setminus (p \cup q)$. By Observation 6, we can either delete or contract (x, y) as a valid d-minor step. Consider the TDAG G'' we get after contraction/deletion. If we delete (x, y) then $G'' = \langle V, E \setminus \{(x, y)\}, s, t\rangle$ contains $p \cup q$, and thus violates our minimality assumption. Similarly if we contract (x, y) where either x or y is not on $p \cup q$. The only problematic case is when (x, y) is contracted (w.l.o.g. forward-contraction) and $\{x, y\} \in p \cup q$. Since forward-contraction is only possible when the outdegree of x is 1, it is impossible that $x \in p$ (as p is a directed path where all nodes except t already have another outgoing edge). Thus $x \in q_{sb}$ or $x \in q_{at}$, and $y \in p \cup q$. Note that $y \neq s$ due to acyclicity.

Case I: $x \in q_{sb}$. We divide into subcases according to where y is on $p \cup q$. The case of $y = t$ will be treated separately. We divide into subcases, showing that in either case $\overline{G}''$ is not series-parallel, which by Lemma 18 means that G'' is not a DSP, i.e., a contradiction. Recall that to show that $\overline{G}''$ is not series-parallel, it is sufficient to show an edge e' and two undirected $s - t$ paths p', q' in $\overline{G}''$ that pass through e' in opposite directions (by conditions (1) and (4) of Prop. 16). Note that in graph G'', nodes x and y are the same node. All subcases are graphically depicted in Figure 5.

1. $y \in q_{sb}$. Rename x, y in G such that x is the one closer to s. We select some $e' \in p_{ab}$, and set $p' = p$ and $q' = q = [q_{sx}] - (x = y) - q_{yb} - [q_{bt}]$.
2. $y \in p_{sa}$. We select some $e' \in p_{ab}$, and set $p' = p = [p_{sx}] - p_{xa} - [p_{ab}] - p_{bt} - t$ and $q' = q$.
3. $y \in q_{at}$. We select some $e' \in q_{xb}$, and set $p' = [p_{sb}] - q_{bx} - (x = y) - q_{yt} - t$ and $q' = [q_{sx}] - q_{xb} - [p_{bt}]$.
4. $y \in p_{ab}$. We select some $e' \in p_{ay}$, and set $p' = p = [p_{sa}] - p_{ay} - (y = x) - p_{yb} - [p_{bt}]$ and $q' = s - q_{sx} - (x = y) - q_{ya} - [q_{at}]$.
5. $y \in \{b\} \cup p_{bt}$. We select some $e' \in p_{ab}$, and set $p' = [p_{sb}] - p_{by} - [p_{yt}]$ and $q' = [q_{sb}] - q_{ba} - [q_{at}]$. Note that $y = b$ is also covered in this case—where $p_{by} = \emptyset$.

6. $y = a$ (thus after contraction $x = y = a$). The subpath q_{xb} starts with an incoming edge to x, which we denote (z, x). Since G'' is a TDAG, (z, x) belongs to some simple directed $s-t$ path $h = s-z-x-t$. Note that $a \notin [h_{sz}]$, as otherwise G'' contains a cycle $[h_{az}] - x - a$. Similarly for any node $a' \in p_{bt}$, thus $[p_{bt}] \cap [h_{sz}] = \emptyset$. We set $p' = [h_{sz}] - (x = a) - p_{ab} - [p_{bt}]$, and p' is a simple undirected $s - t$ path in $\overline{G}''$. We also set $q' = [p_{sa}] - z - q_{zb} - [p_{bt}]$, and note that both paths pass through $e' = (z, a) = (z, x) \in E'$ in opposite directions.

This leaves the subcase of $y = t$. In this case we do not contract (x, t). Instead, we consider the node z that succeeds x on $[q_{sb}]$ (it is possible that $z = b$), and denote $e' = (z, x)$. Set $p' = [p_{sb}] - q_{bz} - z - x - t$ and $q' = [q_{sx}] - z - q_{zb} - [p_{bt}]$, and note that both paths pass through $e' = (z, x)$ in opposite directions. Finally, select and arbitrary edge $e'' \in q_{at}$ and either delete or contract it to get a d-minor G''. We note that in G'' the vertices a, t are still distinct, since if q_{at} is a single edge, this edge can be deleted rather than contracted. The simple undirected paths p', q' still exist in G'' (as neither one uses q_{at}), and thus G'' is not a DSP in contradiction to our minimality assumption.

Case II: $x \in q_{at}$. Again, we divide into subcases. Each subcase is roughly parallel to the subcase with the same index in Case I. We write the proof in full due to slight difference. Note that we need no special treatment for $y = t$.

1. $y \in [q_{at}]$. Rename x, y in G so that y is closer to t (possibly $y = t$ or $x = a$). We select some $e' \in p_{ab}$, and set $p' = p$ and $q' = [q_{sb}] - q_{ba} - a - q_{ax} - (x = y) - q_{yt} - t$.
2. $y \in p_{bt}$. We set $e' = (a, b)$, $p' = p = s - a - b - (y = x) - t$ and $q' = q = s - b - a - (x = y) - t$.
3. $y \in q_{sb}$. We select some $e' \in q_{yb}$, and set $p' = p_{sa} - b - (y = x) - t$, $q' = q_{sy} - b - t$.
4. $y \in p_{ab}$. We select some $e' \in p_{yb}$, and set $p' = p = [p_{sa}] - p_{ay} - (y = x) - p_{yb} - [p_{bt}]$ and $q' = [q_{sb}] - q_{by} - (y = x) - q_{xt} - t$.
5. $y \in \{s\} \cup p_{sa}$. We set $e' = (a, b)$, $p' = p = s - (y = x) - a - b - t$ and $q' = s - b - a - (y = x) - t$.
6. $y = b$. The subpath q_{xt} starts with an edge (z, x), and there is a directed path $h = s - z - x - t$ such that $p_{bt} \cap h_{sz} = \emptyset$ (otherwise the is a cycle). We then set $e' = (z, (x = b))$, $p' = p_{sa} - (b = x) - q_{zt}$ and $q' = h_{sz} - p_{bt}$.

This covers all cases of forward-contraction. The cases of backward-contraction are symmetric. □

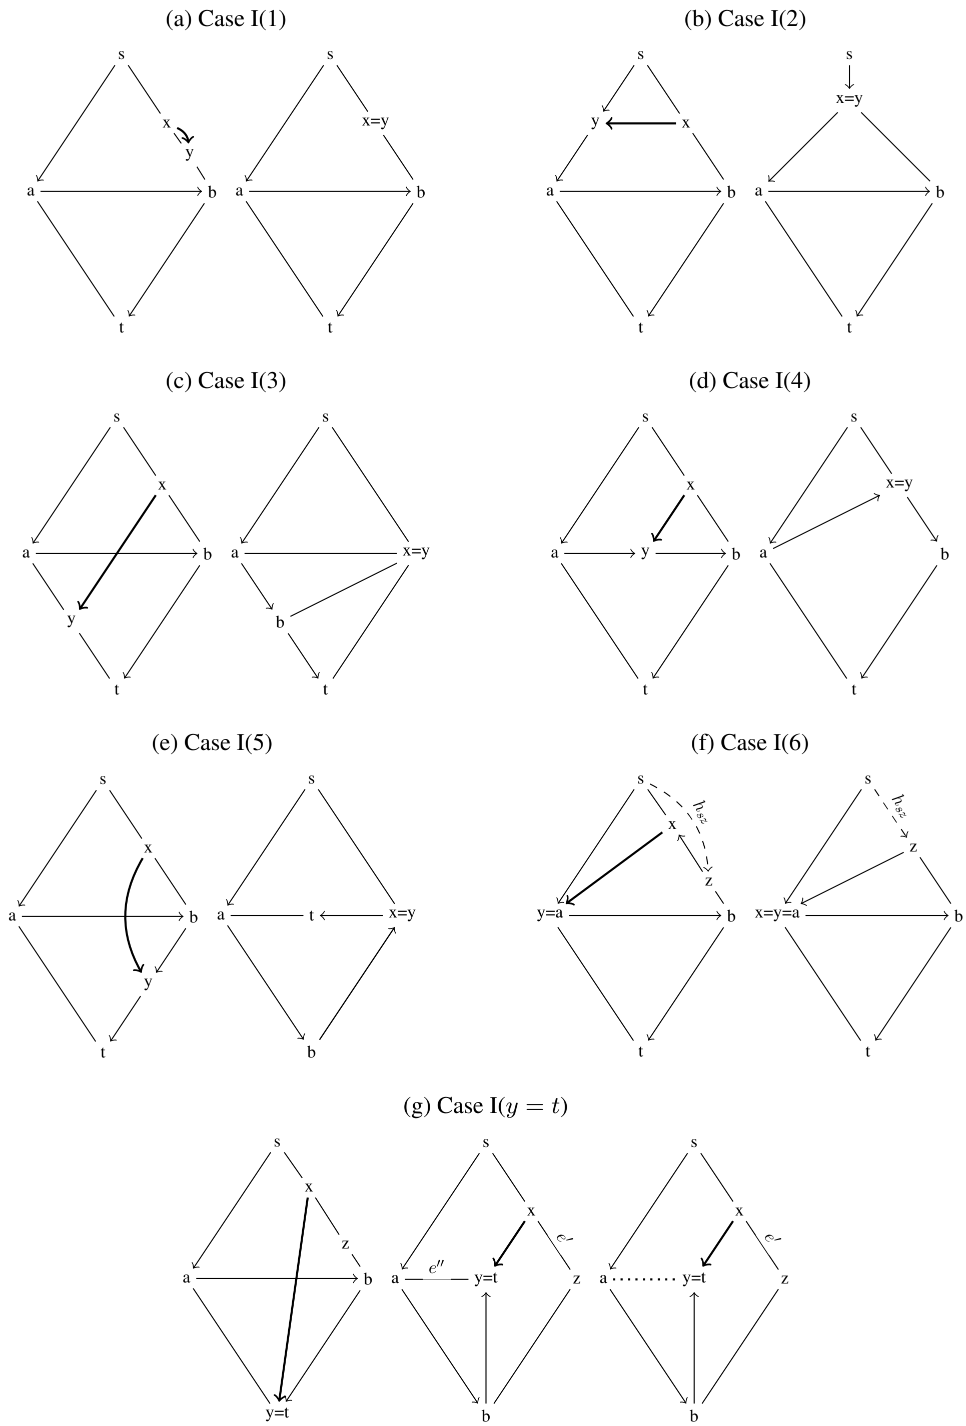

Figure 5: The figure shows all seven subcases of Case I. Each of the first six cases shows a pair of figures, where the figure on the left shows G (or part of it) before contraction, and the figure on the right shows the result after contraction, G''. In the last case we show the original graph G, then the same graph where nodes are rearranged, and in the rightmost figure the dotted line is the path q_{at} after contraction/deletion of e''.