
Enumerating Optimal Cost-Constrained Adjustment Sets

Batya Kenig¹

¹Technion, Israel Institute of Technology, Haifa, Israel

Abstract

Estimating causal effects from observational data is a key problem in causal inference, often addressed through covariate adjustment sets that enable unbiased estimation of interventional means. This paper tackles the challenge of finding optimal covariate adjustment sets under budget constraints, a practical concern in many applications. We present algorithms for enumerating valid and minimal adjustment sets up to a specified cost, ordered by their proximity to outcome variables, which coincides with estimator variance. Our approach builds on existing graphical criteria and extends them to accommodate budgetary considerations, providing a useful tool for addressing resource limitations.

1 INTRODUCTION

Estimating the causal effect between treatment and outcome variables in the presence of, possibly hidden, confounding variables is a fundamental problem in causal inference [Pearl, 2009]. When causal-effect estimation is based on observational data, confounders pose a major problem because they act as hidden influencers that impact both the cause and effect, creating a biased association. The *interventional mean* measures the expected outcome under a specific intervention, serving as the basis for other causal effect measures. Estimating the interventional mean requires computing the post-interventional distribution $P(Y|do(\mathbf{X}))$, reflecting the distribution of outcome Y after intervening on treatment variables X . *Covariate adjustment sets* are sets of variables that enable computing an unbiased estimate of the interventional mean. An adjustment set is *valid* if it enables unbiased estimation of the causal effect from the joint distribution over the observed variables.

Graphical criteria for selecting adjustment sets have been extensively studied. Pearl’s back-door criterion [1993] is well-

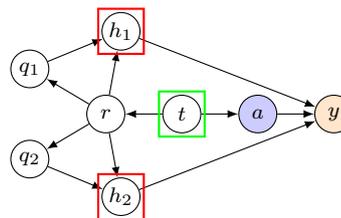


Figure 1: Both $Z_1 = \{h_1, h_2\}$ and $Z_2 = \{t\}$ are valid adjustment sets for estimating the causal effect of a on y , but Z_1 yields an estimator with lower variance when compared to the estimator generated using Z_2 (note that h_1, h_2 are adjacent to y). On the other hand, $|Z_2| < |Z_1|$, and hence potentially cheaper to measure, and apply for adjustment. Example based on Rotnitzky [2021].

known but incomplete; that is, it cannot be used to identify all valid adjustment sets. Recent advancements [Perkovic et al., 2017, van der Zander et al., 2019] have provided sound and complete criteria for various types of graphs, including those with unobserved variables.

When based on valid adjustment sets, estimators of causal effects are unbiased, but their variances can differ significantly across different adjustment sets [Smucler and Rotnitzky, 2022, Runge, 2021]. This has led to extensive research on identifying adjustment sets that yield estimators with minimal asymptotic variance. Rotnitzky and Smucler [2020] derived a graphical characterization of the optimal adjustment set in non-parametric models. Smucler et al. [2021] provided graphical criteria for optimal *minimal*, and *minimum cardinality* adjustment sets in causal models with hidden variables. An adjustment set is *minimal* if it does not contain any adjustment set as a proper subset, and is *minimum cardinality* if no smaller adjustment set exists. The motivation for minimum-cardinality adjustment sets stems from the fact that computing the causal effect estimator requires summing over all possible values of the adjustment set, leading to an exponential dependence on its domain size. This computational burden grows rapidly as the number of covariates and their possible values increase.

In realistic settings, choosing an adjustment set involves balancing precision and cost. For example, variables requiring laboratory tests may be much more expensive to measure than variables pertaining to clinical examination. Such cost considerations motivate the problem of finding adjustment sets whose overall cost meets a given budget constraint [Rotnitzky and Smucler, 2020, Smucler and Rotnitzky, 2022]. The basic budget constraint is the one that places a limit k on the size of the adjustment set. Allowing integral weights on the vertices of the model grants us wider flexibility in selecting adjustment sets. In particular, the weight of an adjustment set can account for not only its cardinality, but also its domain size, and measurement-cost, which directly impact the efficiency of computing the causal effect.

Recent work [Smucler et al., 2021, Smucler and Rotnitzky, 2022, Runge, 2021] has established that valid adjustment sets that are *closer* to the outcome variables Y yield estimators with smaller asymptotic variance for all distributions that factorize according to the causal DAG; we later make this notion precise. For example, in Figure 1, we present a causal model from Rotnitzky [2021], where both $Z_1 = \{h_1, h_2\}$ and $Z_2 = \{t\}$ are valid adjustment sets for measuring the causal effect of a on y . Since Z_1 is *closer* to y than Z_2 , then Z_1 will yield a superior estimator in terms of variance when compared to Z_2 , while Z_2 is superior in terms of cost (i.e., $|Z_2| < |Z_1|$). This example illustrates the tension between the precision of the estimator and its cost.

In causal inference, selecting an adjustment set is a multi-criteria optimization problem: different sets trade off accuracy, as measured by the variance of the causal-effect estimator, against cost, which reflects the expense of measuring the covariates, and the computational cost they yield for computing the estimator. Crucially, there is no single optimal solution when these criteria compete (i.e., what improves one may worsen the other). A natural approach is to enumerate the *Pareto frontier* of adjustment sets, which is the family of valid adjustment sets that are undominated with respect to these competing objectives. We show that this frontier corresponds exactly to the class of *important separators* Marx [2011], Cygan et al. [2015] in a certain undirected graph derived from the causal model. Applying the concept of important separators to the task of finding optimal adjustment sets yields a principled and efficient algorithmic solution for enumerating all Pareto-optimal adjustment sets of size at most k .

However, the approach relying on important separators of size at most k inherently targets small adjustment sets because its runtime depends exponentially on k , and is therefore applicable only when the optimal sets are small (see Theorem 1). To go beyond this regime, we develop an efficient, general algorithm for ranked enumeration of all valid adjustment sets, ordered first by cost and then by variance that achieves *polynomial-delay* regardless of the size of the adjustment sets returned. This broader enumeration enables

exploring the trade-offs even when optimal sets are large.

Finally, we present an algorithm that goes beyond the enumeration of Pareto-optimal adjustment sets to generate all minimal, valid adjustment sets of size at most k , ranked by their vicinity to the outcome variable, which is a proxy for estimator variance. While important separators capture all Pareto-optimal solutions with respect to cost and variance, this subset may be very small when compared to the entire set of minimal adjustment sets of size k , limiting flexibility in practice. Moreover, assigning precise costs to covariates is often difficult: different variables may be assigned the same weight despite differences in availability, reliability, or measurement burden. This has been observed, for example, in a biomolecular causal study Taheri et al. [2023] where a cheaper protein (PI3K) yielded comparable precision to a costlier one (Ras), even though both were treated equally under the cost model. In such cases, enumerating near-optimal adjustment sets and not just Pareto-optimal ones, may potentially uncover practically preferable options that standard optimization may overlook.

Contributions. We assign each variable in the causal model an integral weight representing the cost of measuring it and including it in an adjustment set. Unweighted models correspond to those where all variable weights are simply 1. A key contribution of this work is the characterization of all Pareto-optimal adjustment sets of weight at most k . An adjustment set is Pareto-optimal if it is valid, and every other valid adjustment set either has higher cost or higher estimator variance in all distributions consistent with the causal DAG. We show that these sets correspond exactly to the class of important separators in a certain undirected graph derived from the causal DAG, allowing us to efficiently enumerate the entire Pareto frontier using tools from the theory of parameterized algorithms.

Theorem 1. Let G be a causal DAG, with an integral weight function $w : \mathcal{V}(G) \rightarrow \{1, \dots, c\}$ for some constant c , and $X, Y \subseteq \mathcal{V}(G)$ be disjoint. There is an algorithm that lists the Pareto-Optimal minimal adjustment sets in G , for computing an unbiased estimator of the interventional mean of the outcomes Y under interventions on X , of size at most k , in time $O(4^k \cdot k \cdot (n+m))$ where $n = |\mathcal{V}(G)|$, and $m = |\mathcal{E}(G)|$.

Moving beyond the Pareto-Optimal frontier, we present an algorithm that returns all valid adjustment sets, minimal and non-minimal, ranked by cost and, secondarily, by their closeness to the outcome variables. This algorithm generalizes the result of Smucler et al. [2021], which identified the minimum-cardinality adjustment set that is closest to the outcome. Our method supports fully ranked exploration of the tradeoff space under integral cost functions, even without a fixed bound on the cost.

Theorem 2. Let G be a causal DAG, with an integral weight

function $w : \mathcal{V}(G) \rightarrow \mathbb{N}_{\geq 1}$, and let $X, Y \subseteq \mathcal{V}(G)$ be disjoint. There exists an enumeration algorithm that outputs all valid adjustment sets in G for computing an unbiased estimator of the interventional mean of the outcomes Y under interventions on X . The adjustment sets are listed in order of non-decreasing total weight, and ties are broken by proximity to the outcome variables Y . The delay of the algorithm is $O(Kn \cdot T(n, m))$, where K is the size of the largest adjustment set listed, $n = |\mathcal{V}(G)|$, $m = |\mathcal{E}(G)|$, and $T(n, m)$ denotes the time to compute a minimum separator in an undirected graph with n vertices and m edges.

Finding a minimum separator can be reduced, by standard techniques [Even and Even, 2012], to the problem of finding a maximum flow in the graph [Ford and Fulkerson, 2010]. Currently, the fastest known algorithm for max-flow runs in almost linear time $m^{1+o(1)}$ [Chen et al., 2022]. In the rest of this paper, we denote by $T(n, m)$ the time to find a minimum s, t -separator in an undirected graph.

Theorem 1 presents an algorithm for enumerating the subset of Pareto-optimal minimal adjustment sets, while Theorem 2 provides a polynomial-delay algorithm for enumerating all valid adjustment sets, including non-minimal ones, ranked by cost and proximity to the outcome variables. Ideally, we would like an efficient, polynomial-delay algorithm that enumerates only minimal adjustment sets, ranked by cost with a secondary ranking by proximity to the outcome, which serves as a proxy for estimator variance. However, we later present a hardness result that rules out such an algorithm using known techniques. This suggests that achieving efficient enumeration with both minimality and ranking guarantees may require fundamentally new algorithmic ideas. To bridge the gap between these two settings, namely enumerating only minimal adjustment sets, but going beyond the Pareto frontier, we develop an *FPT-delay* algorithm that lists all minimal valid adjustment sets of cost at most k , ranked by their proximity to the outcome. This addresses scenarios where important separators are too few or too restrictive, allowing practitioners to explore a broader space of high-quality adjustment sets efficiently.

Theorem 3. Let G be a causal DAG, with an integral weight function $w : \mathcal{V}(G) \rightarrow \{1, \dots, c\}$ for some constant c , and $X, Y \subseteq \mathcal{V}(G)$ be disjoint. There is an algorithm that lists all minimal adjustment sets in G for computing an unbiased estimator of the interventional mean of the outcomes Y under interventions on X , of weight at most k , with delay $O(k^2 4^k (n+m))$ where $n=|\mathcal{V}(G)|$, and $m=|\mathcal{E}(G)|$. The algorithm outputs the minimal separators in non-decreasing order of their distance from Y .

Theorems 1 and 3 are stated for the case where the weight function assigns each variable an integer in $\{1, \dots, c\}$ for some constant c . In the proofs, we assume all weights are 1, and extend the results to bounded weights using the standard vertex-splitting technique. Specifically, each vertex of

weight $w(v) \leq c$ is replaced by $w(v)$ unweighted vertices forming a clique, preserving separation structure and total weight. This transformation yields an equivalent unweighted graph where our algorithms can be applied directly. For complete technical details, see Section B in the Appendix.

In our results, we leverage the work of van der Zander et al. [2014, 2019], translating the problem into one of finding *separators* and *minimal separators* in an undirected graph that is derived from the so-called *proper backdoor graph*. Let $s, t \in \mathcal{V}(G)$ be two distinguished vertices in a finite, simple, undirected graph $G(\mathcal{V}, \mathcal{E})$. An s, t -separator is a subset $S \subseteq \mathcal{V}(G)$, such that removing S and its incident edges disconnects s and t in G . An s, t -separator $S \subseteq \mathcal{V}(G)$ is a minimal s, t -separator if no strict subset of S is also an s, t -separator. A key technical ingredient in our approach is a simple yet powerful graphical criterion for comparing the asymptotic variance of estimators associated with different adjustment sets. This criterion operates directly on the structure of minimal separators in the derived undirected graph. Beyond their implications for causal inference, our algorithms also contribute to the study of ranked enumeration of separators in graphs.

Previous work on separator enumeration. Enumerating minimal separators of bounded cardinality (or weight) refines and extends two well-studied enumeration problems: enumeration of all minimal separators, and enumeration of all minimum-cardinality separators [Kanevsky, 1990]. Berry et al. [2000] developed an efficient algorithm that lists the minimal separators of an undirected graph \mathcal{H} with a delay of $O(|\mathcal{V}(\mathcal{H})|^3)$ between consecutive outputs. The algorithm of Berry et al. [2000], as well as others [Kloks and Kratsch, 1998, Takata, 2010, Shen and Liang, 1997], does not list the minimal separators in any ranked order, and cannot restrict the output only to separators of weight at most k . Kanevsky [1990] developed a complicated algorithm that enumerates all the minimum-cardinality separators of a graph; Theorem 2 strictly generalizes this result.

Challenges and techniques. In a ranked enumeration algorithm, finding the top element is basically an *optimization problem*. In our case, there are well known algorithms for finding a *minimum-weight s, t -separator* [Henzinger et al., 2000, Even and Tarjan, 1975, Chen et al., 2022]. For $K > 1$, finding the K -th ranking item amounts to computing the optimal minimal s, t -separator under the restriction that it is not among the first $K - 1$ items previously returned. Handling this constraint is the main challenge when designing ranked enumeration algorithms.

The technique of [Lawler et al., 1980] provides a general framework for ranked enumeration corresponding to discrete optimization problems. The main idea is to reduce a ranked enumeration problem to an optimization problem with constraints [Golenberg et al., 2011]. In the standard

approach to applying the Lawler-Murty technique, the algorithm first finds the optimal solution S (e.g., minimum s, t -separator). Then, the subspace of solutions (excluding S) is partitioned using *inclusion* and *exclusion* constraints. The straightforward approach to applying the Lawler-Murty method to ranked enumeration of minimal s, t -separators is by solving the following optimization problem: find the minimum-weight, minimal s, t -separator in the graph G that excludes a subset $U \subseteq \mathbb{V}(G)$, and includes a subset $I \subseteq \mathbb{V}(G)$ of vertices. Using this approach, we immediately hit an obstacle. In Section D of the Appendix we show that deciding whether there exists a minimal s, t -separator that includes a distinguished vertex $v \in \mathbb{V}(G)$, is NP-complete by reduction from the 3-IN-A-PATH problem [Derhy and Picouleau, 2009]. For this reason, our algorithm lists the minimal s, t -separators of weight at most k in *FPT-delay* [Creignou et al., 2017] with parameter k . Here as well, our approach makes use of the notion of important separators [Marx, 2011].

Organization. In Section 2, we provide background on separators in undirected graphs. Section 3 covers causal graphical models and adjustment sets, and reviews the results of [van der Zander et al., 2019, Smucler et al., 2021] that allow translating the problem of finding adjustment sets in causal models to that of finding separators in an undirected graph. In Section 4, we show how the result of Smucler et al. [2021] for comparing two adjustment sets based on the asymptotic variance of their estimator translates to a simple graphical criterion for comparing separators, and we apply the concept of important separators to prove Theorem 1. Section 5 considers the enumeration of all separators in ranked order by weight, where the secondary ranking corresponds to the quality of the associated adjustment sets (Theorem 2). Due to space restrictions, some of the proofs are deferred to the Appendix. The proof of Theorem 3, which presents an algorithm for enumerating minimal separators of bounded weight, ranked by the quality of the corresponding adjustment sets, appears in Section B of the Appendix.

2 BACKGROUND: UNDIRECTED GRAPHS AND SEPARATORS

Let G be an undirected graph with nodes $\mathbb{V}(G)$ and edges $\mathbb{E}(G)$, where $n = |\mathbb{V}(G)|$, and $m = |\mathbb{E}(G)|$. A strictly positive, integral weight function $w : \mathbb{V}(G) \rightarrow \mathbb{N}_{\geq 1}$ is defined on the vertices. For unweighted graphs, we assume $w(v) = 1$ for all $v \in \mathbb{V}(G)$. For a subset of vertices $S \subseteq \mathbb{V}(G)$, the weight of S is $w(S) \stackrel{\text{def}}{=} \sum_{v \in S} w(v)$. For $A, B \subseteq \mathbb{V}(G)$, we abbreviate $AB \stackrel{\text{def}}{=} A \cup B$; for $v \in \mathbb{V}(G)$ we abbreviate $vA \stackrel{\text{def}}{=} \{v\} \cup A$. Let $v \in \mathbb{V}$. We denote by $N_G(v) \stackrel{\text{def}}{=} \{u \in \mathbb{V}(G) : (u, v) \in \mathbb{E}(G)\}$ the neighborhood of v , and by $N_G[v] \stackrel{\text{def}}{=} N_G(v) \cup \{v\}$ the closed

neighborhood of v . For a subset of vertices $T \subseteq \mathbb{V}(G)$, we denote by $N_G(T) \stackrel{\text{def}}{=} \bigcup_{v \in T} N_G(v) \setminus T$, and $N_G[T] \stackrel{\text{def}}{=} N_G(T) \cup T$. We denote by $G[T]$ the subgraph of G induced by T . Formally, $\mathbb{V}(G[T]) = T$, and $\mathbb{E}(G[T]) = \{(u, v) \in \mathbb{E}(G) : \{u, v\} \subseteq T\}$. For a subset $S \subseteq \mathbb{V}(G)$, we abbreviate $G-S \stackrel{\text{def}}{=} G[\mathbb{V}(G) \setminus S]$; for $v \in \mathbb{V}(G)$, we abbreviate $G-v \stackrel{\text{def}}{=} G-\{v\}$. We say that G' is a *subgraph* of G if it results from G by removing vertices and edges; formally, $\mathbb{V}(G') \subseteq \mathbb{V}(G)$ and $\mathbb{E}(G') \subseteq \mathbb{E}(G)$.

Let $e = (u, v) \in \mathbb{E}(G)$. The *contraction* of e results in a new graph G' , where u and v are identified with a new vertex w_e that is adjacent to $N_G(u) \cup N_G(v)$. Formally, $\mathbb{V}(G') = \mathbb{V}(G) \setminus \{u, v\} \cup \{w_e\}$, and $\mathbb{E}(G') = \mathbb{E}(G) \setminus \{e\} \cup \{(w_e, y) : y \in N_G(u) \cup N_G(v)\}$. The *contraction* of e to vertex u results in the graph G' , where v is identified with u that is adjacent to $N_G(u) \cup N_G(v)$. Formally, $\mathbb{V}(G') = \mathbb{V}(G) \setminus \{v\}$, and $\mathbb{E}(G') = \mathbb{E}(G) \setminus \{e\} \cup \{(u, y) : y \in N_G(v)\}$.

Let $u, v \in \mathbb{V}(G)$. A *simple path* between u and v , called a u, v -path, is a finite sequence of distinct vertices $u = v_1, \dots, v_k = v$ where, for all $i \in [1, k-1]$, $(v_i, v_{i+1}) \in \mathbb{E}(G)$, and whose ends are u and v . A u, v -path is *chordless* or *induced* if $(v_i, v_j) \notin \mathbb{E}(G)$ whenever $|i - j| > 1$.

We say that a subset of vertices $V' \subseteq \mathbb{V}(G)$ is *connected* in G if $G[V']$ contains a path between every pair of vertices in V' . A subset of vertices $V' \subseteq \mathbb{V}(G)$ is called a *connected component* of G if V' is connected, and $G[V' \cup \{x\}]$ is not connected for every $x \in \mathbb{V} \setminus V'$. We say that G is connected if $\mathbb{V}(G)$ is connected. Let $V_1, V_2 \subseteq \mathbb{V}(G)$ denote two disjoint vertex subsets of $\mathbb{V}(G)$. We say that V_1 and V_2 are adjacent if there is at least one pair of adjacent vertices $v_1 \in V_1$ and $v_2 \in V_2$. We say that there is a path between V_1 and V_2 if there exist vertices $v_1 \in V_1$ and $v_2 \in V_2$ such that there is a path between v_1 and v_2 .

Let $u \in \mathbb{V}(G)$; we denote by $\text{Sat}(G, u)$ the graph that results from G by adding edges between all pairs of vertices in $N_G[u]$. In other words, $\text{Sat}(G, u)$ is the graph where the set $N_G[u]$ has been *saturated*, and forms a clique. Formally, $\mathbb{V}(\text{Sat}(G, u)) \stackrel{\text{def}}{=} \mathbb{V}(G)$, and $\mathbb{E}(\text{Sat}(G, u)) \stackrel{\text{def}}{=} \mathbb{E}(G) \cup \{(x, y) : x, y \in N_G[u]\}$. For a set of vertices $U \subseteq \mathbb{V}(G)$, we denote by $\text{Sat}(G, U)$ the graph that results by adding edges between all vertices in $N_G[u]$ for all $u \in U$. Formally, $\mathbb{V}(\text{Sat}(G, U)) = \mathbb{V}(G)$ and

$$\mathbb{E}(\text{Sat}(G, U)) \stackrel{\text{def}}{=} \mathbb{E}(G) \cup \bigcup_{u \in U} \{(x, y) : x, y \in N_G[u]\}. \quad (1)$$

2.1 MINIMAL SEPARATORS

Let $s, t \in \mathbb{V}(G)$. For $X \subseteq \mathbb{V}(G)$, we let $\mathcal{C}(G-X)$ denote the set of connected components of $G-X$. The vertex set X is called a *separator* of G if $|\mathcal{C}(G-X)| \geq 2$, an s, t -separator if s and t are in different connected

components of $\mathcal{C}(G-X)$, and a *minimal s, t -separator* if no proper subset of X is an s, t -separator of G . For an s, t -separator X , we denote by $C_s(G-X)$ and $C_t(G-X)$ the connected components of $\mathcal{C}(G-X)$ containing s and t respectively. In other words, $C_s(G-X) = \{v \in V(G) : \text{there is a path from } s \text{ to } v \text{ in } G-X\}$.

Lemma 1. [Berry et al., 2000] An s, t -separator $X \subseteq V(G)$ is a minimal s, t -separator if and only if $N_G(C_s(G-X)) = N_G(C_t(G-X)) = X$.

A subset $X \subseteq V(G)$ is a *minimal separator* if there exist a pair of vertices $u, v \in V(G)$ such that X is a minimal u, v -separator. A connected component $C \in \mathcal{C}(G-X)$ is called a *full component* of X if $N_G(C) = X$. By Lemma 1, X is a minimal u, v -separator if and only if the components $C_u(G-X)$ and $C_v(G-X)$ are full components of X . We denote by $\mathcal{S}_{s,t}(G)$ the set of minimal s, t -separators of G , and by $\mathcal{S}(G)$ the set of minimal separators of G .

An immediate consequence of Lemma 1 is the following, which will be used later on. Proof deferred to Appendix A.

Proposition 1. Let $S_1, S_2 \in \mathcal{S}_{s,t}(G)$. Then $C_s(G-S_1) \subseteq C_s(G-S_2)$ if and only if $C_t(G-S_2) \subseteq C_t(G-S_1)$.

2.1.1 Separators Between Vertex-Sets.

Let $A, B \subseteq V(G)$ be disjoint and non-adjacent. A subset $S \subseteq V(G) \setminus AB$ is an A, B -separator if, in the graph $G-S$, there is no path between A and B . We say that S is a minimal A, B -separator if no proper subset of S is an A, B -separator. We denote by $\mathcal{S}_{A,B}(G)$ the set of minimal A, B -separators of G . In Section A.1 of the Appendix, we show how separators between vertex-sets can be represented as separators between singleton vertices.

Theorem 4. Let $A, B \subseteq V(G)$ be disjoint and non-adjacent, where $s \in A$ and $t \in B$. Let H be the graph that results from G by: (1) adding all edges between s and $N_G(A)$, (2) adding all edges between t and $N_G(B)$, and (3) removing vertices $AB \setminus \{s, t\}$ and their adjacent edges. Then $\mathcal{S}_{s,t}(H) = \mathcal{S}_{A,B}(G)$.

2.1.2 Close Separators.

When $S \in \mathcal{S}_{s,t}(G)$ where $S \subseteq N_G(s)$, then we say that S is *close to s* [Kloks and Kratsch, 1998].

Lemma 2. [Kloks and Kratsch, 1998] If s and t are non-adjacent, then there exists exactly one minimal s, t -separator that is close to s .

Let $A, B \subseteq V(G)$ be disjoint and non-adjacent. From Lemma 2, and Theorem 4, we get that there exists a unique minimal A, B -separator that is close to A . If $a \in A, b \in B$, and H is the graph that results from G by adding all edges

between a and $N_G(A)$ and all edges between b and $N_G(B)$, and then removing $AB \setminus \{a, b\}$ and their adjacent edges, we get by Theorem 4 that $\mathcal{S}_{a,b}(H) = \mathcal{S}_{A,B}(G)$. By Lemma 2 there is unique minimal a, b -separator $S \in \mathcal{S}_{a,b}(H)$ that is close to a , where $S \subseteq N_H(a) \subseteq N_G(A)$. Due to the equivalence $\mathcal{S}_{a,b}(H) = \mathcal{S}_{A,B}(G)$, we get that S is the unique minimal A, B -separator that is closest to A .

Definition 1 (closer to, \preceq). Let $S_1, S_2 \in \mathcal{S}_{s,t}(G)$. We say that S_1 is *strictly closer to s* than S_2 , denoted $S_1 \prec S_2$ if $C_s(G-S_1) \subset C_s(G-S_2)$, and that S_1 is *closer to s* than S_2 , denoted $S_1 \preceq S_2$ if $C_s(G-S_1) \subseteq C_s(G-S_2)$.

2.2 MINIMUM SEPARATORS

A subset $S \subseteq V(G)$ is a *minimum-weight s, t -separator* of G (or just minimum s, t separator) if $w(S) \leq w(S')$ for every other s, t -separator S' . We denote by $\kappa_{s,t}(G)$ the weight of a minimum s, t -separator of G , and by $\mathcal{L}_{s,t}(G)$ the set of all minimum s, t -separators of G . Finding a minimum s, t -separator can be reduced, by standard techniques [Even and Even, 2012], to the problem of finding a maximum flow in the graph [Ford and Fulkerson, 2010]. Currently, the fastest known algorithm for max-flow runs in almost linear time $m^{1+o(1)}$ [Chen et al., 2022]. We denote by $T(n, m)$ the time to find a minimum s, t -separator of G .

The following theorem is a straightforward extension of a known result for unweighted graphs Cygan et al. [2015], adapted here to the weighted setting. Its proof is deferred to Section A.2 of the Appendix.

Theorem 5. There exists a unique minimum s, t -separator $S^* \in \mathcal{L}_{s,t}(G)$ such that $S^* \preceq S$ for all $S \in \mathcal{L}_{s,t}(G)$, and S^* can be found in time $O(n \cdot T(n, m))$.

2.3 IMPORTANT MINIMAL SEPARATORS

The notion of *important separators* has been applied to the design of various fixed-parameter tractable algorithms [Cygan et al., 2015, Marx, 2011].

Definition 2. [Cygan et al., 2015, Marx, 2011] Let $S \subseteq V(G)$. We say that S is an *important s, t -separator* if $S \in \mathcal{S}_{s,t}(G)$, and for any other $S' \in \mathcal{S}_{s,t}(G)$ it holds that:

$$C_s(G-S') \subset C_s(G-S) \implies |S'| > |S|.$$

In what follows, we denote by $\mathcal{S}_{s,t}^*(G)$ the set of important s, t -separators, and by $\mathcal{S}_{s,t,k}^*(G)$ the set of important s, t -separators whose size is at most k .

Theorem 6. Cygan et al. [2015], Marx [2011] There are at most 4^k important s, t -separators of G whose size is at most k , and there is an algorithm that outputs them in total time $O(4^k \cdot k \cdot (n + m))$.

3 BACKGROUND: CAUSAL GRAPHICAL MODELS

Let G be a directed graph with nodes $V(G)$ and edges $E(G)$. As in the undirected case, given a set of vertices $Z \subseteq V(G)$, we denote by $G[Z]$ the directed subgraph induced by Z . Formally, $E(G[Z]) = \{(u, v) \in E(G) : \{u, v\} \subseteq Z\}$. A *path* between u and v is a sequence of adjacent vertices (v_1, \dots, v_j) such that $u = v_1$ and $v = v_j$. A vertex w is a *collider* on a path if the path contains the subpath $u \rightarrow w \leftarrow v$. The path is directed, or *causal*, if $v_i \rightarrow v_{i+1}$ for all $i \in \{1, 2, \dots, j-1\}$. A directed cycle is a directed path from u to v , combined with the directed edge $v \rightarrow u$. A directed acyclic graph (DAG) is a directed graph without directed cycles. The *moral graph* of a DAG G is an undirected graph G^m with the same vertex set as G , and where $(u, v) \in E(G^m)$ if and only if there is a directed edge between u and v in G , or if there exists a vertex $w \in V(G)$ such that $u \rightarrow w \leftarrow v$ is an induced subgraph of G .

If $u \rightarrow w \in E(G)$, then u is a *parent* of w . If there is a directed path from u to w , then u is an *ancestor* of w and w a *descendant* of u . We follow the convention that a vertex is an ancestor and descendant of itself. The parents, ancestors and descendants of $w \in V(G)$ are denoted by $\text{pa}_G(w)$, $\text{an}_G(w)$ and $\text{de}_G(w)$, respectively. For $Z \subseteq V(G)$, we define $\text{an}_G(Z) \stackrel{\text{def}}{=} \bigcup_{z \in Z} \text{an}_G(z)$, $\text{de}_G(Z) \stackrel{\text{def}}{=} \bigcup_{z \in Z} \text{de}_G(z)$, and $\text{pa}_G(Z) \stackrel{\text{def}}{=} \bigcup_{z \in Z} \text{pa}_G(z)$. We denote by $\text{nd}_G(Z) \stackrel{\text{def}}{=} V(G) \setminus \text{de}_G(Z)$ the nondescendants of Z .

Let G be a DAG, and $x, y \in V(G)$. The set of *causal vertices* between x and y , denoted $\text{cv}_G(x, y)$, are those vertices that lie on a directed x, y -path in G . The set of *forbidden vertices* in G with respect to $x, y \in V(G)$ is defined to be $\text{forb}_G(x, y) \stackrel{\text{def}}{=} \{x\} \cup \text{de}_G(\text{cv}_G(x, y))$. Accordingly, for $X, Y \subseteq V(G)$, we define $\text{cv}_G(X, Y)$ as the vertices that lie on a directed x, y -path in G for any $x \in X$ and $y \in Y$. Given $X, Y \subseteq V(G)$, the *proper back-door graph* $G^{\text{pbd}}(X, Y)$ [van der Zander et al., 2019] is defined as the graph that results from G by removing from G the first edge of every directed X, Y -path. Formally,

$$E(G^{\text{pbd}}(X, Y)) \stackrel{\text{def}}{=} E(G) \setminus \{x \rightarrow u : x \in X, u \in \text{cv}_G(X, Y)\} \quad (2)$$

Identification via Covariate Adjustment

A *Bayesian Network* (BN) for a set of variables $V = \{v_1, \dots, v_n\}$ is a pair $\mathcal{B} \stackrel{\text{def}}{=} (G, P)$ where G is a DAG, and P a joint probability distribution for V that factorizes as $P(V) = \prod_{i=1}^n P(v_i | \text{pa}_G(v_i))$. For a variable $u \in V$, we denote by \mathbf{u} an assignment to u , and for a subset of variables $X \subseteq V$, we denote by \mathbf{X} an assignment to all variables in X . We say that the BN \mathcal{B} is *causal* if every edge $v_i \rightarrow v_j \in E(G)$ represents a direct causal effect of v_i on v_j . Given a causal BN $\mathcal{B} = (G, P)$ and a subset $X \subseteq V$, the

post intervention distribution in X is:

$$P(\mathbf{V} | do(\mathbf{X})) \stackrel{\text{def}}{=} \begin{cases} \prod_{v_i \in V \setminus X} P(v_i | \text{pa}_G(v_i)) & \mathbf{V} \text{ is consistent with } \mathbf{X} \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

where \mathbf{V} is consistent with \mathbf{X} if \mathbf{V} and \mathbf{X} assign the same values to the variables in $X \cap V$, and $do(\mathbf{X})$ represents an intervention that sets $X = \mathbf{X}$. In a DAG, this intervention corresponds to removing all edges into X (i.e., all edges between X and $\text{pa}_G(X)$). The term $P(Y | do(X = \mathbf{X}))$ is called the *causal effect* of X on Y .

When all variables $V(G)$ in the DAG G are observed, the causal effect $P(Y | do(X = \mathbf{X}))$ of X on Y in G can be determined directly from P using (3). When some variables are unobserved, then $P(Y | do(X = \mathbf{X}))$ cannot necessarily be computed directly from P . When it can, then it is said that the causal effect of X on Y is *identifiable* [Pearl, 2009].

Definition 3. [Pearl, 2009] Given a DAG G , and pairwise disjoint $X, Y, Z \subseteq V(G)$, Z is called an *adjustment* for estimating the causal effect of X on Y if, for every distribution P that factorizes according to G , it holds:

$$P(\mathbf{Y} | do(\mathbf{X})) = \begin{cases} P(\mathbf{Y} | \mathbf{X}) & \text{if } Z = \emptyset \\ \sum_{\mathbf{Z}} P(\mathbf{Y} | \mathbf{X}, \mathbf{Z}) P(\mathbf{Z}) & \text{otherwise} \end{cases} \quad (4)$$

Let $R \subseteq V(G)$ denote the set of observable variables in $V(G)$. We say that an adjustment set Z is *valid* if $Z \subseteq R$. There exist graphs for which $P(Y | do(X = \mathbf{X}))$ is identifiable but for which no valid adjustment set $Z \subseteq R$ exists. Smucler et al. [2021], and van der Zander et al. [2019] established a graphical criterion that determines whether $P(Y | do(X = \mathbf{X}))$ has a valid adjustment set.

Let $I \subseteq R \subseteq V(G)$. We denote by $\mathcal{A}_{X,Y}(I, R, G)$ all of the adjustment sets Z for X, Y in G according to (4), where $I \subseteq Z \subseteq R$. We call the set $\mathcal{A}_{X,Y}(I, R, G)$ the I, R adjustment sets for X, Y in G . We say that Z is a *minimal I, R adjustment set* for X, Y if no proper subset of Z is an I, R adjustment set for X, Y . We denote by $\mathcal{A}_{X,Y}^{\min}(I, R, G)$ the minimal I, R adjustment sets for X, Y in G .

3.1 CHARACTERIZING ADJUSTMENT SETS AS SEPARATORS IN AN UNDIRECTED GRAPH

In this section, we consider BNs with DAG G , where $R \subseteq V(G)$ is the set of observable variables. We assume that $\emptyset \subset R$ since otherwise, no valid adjustment set exists. Next, we describe the results of van der Zander et al. [2019], and Smucler et al. [2021], which showed that $\mathcal{A}_{X,Y}(I, R, G)$, the I, R adjustment sets for X, Y in G , are represented as separators in a certain undirected graph that we describe next. Given a DAG G , two distinguished vertex-sets $X, Y \subseteq V(G)$, a subset of observable variables $R \subseteq V(G)$, and a

subset $I \subseteq R$, we define the undirected graph:

$$\mathcal{H}_{X,Y}^0(I,G) \stackrel{\text{def}}{=} \left(G^{\text{pbd}}(X,Y)[\text{an}_G(I \cup X \cup Y)] \right)^m \quad (5)$$

where G^{pbd} is the proper back-door graph (see (2)), and $(\dots)^m$ refers to the undirected moral graph.

Definition 4 ($\mathcal{H}_{X,Y}^1(I,G)$). $\mathcal{H}_{X,Y}^1(I,G)$ is the undirected graph that results from $\mathcal{H}_{X,Y}^0(I,G)$ by: (1) adding all edges between $X \cup Y$ and I , and (2) saturating (see (1)) all vertices of $\mathcal{H}_{X,Y}^0(I,G)$ that belong to $(\text{V}(G) \setminus R) \cup \text{f}\circ\text{r}\text{b}_G(X,Y)$, and removing them from the graph.

Following, is the key result that relates (X,Y) -separators in $\mathcal{H}_{X,Y}^1(I,G)$ and the valid adjustment sets $\mathcal{A}_{X,Y}(I,R,G)$.

Theorem 7. [van der Zander et al., 2019, Smucler et al., 2021]

1. $\mathcal{A}_{X,Y}(I,R,G) \neq \emptyset$ if and only if X and Y are not adjacent in $\mathcal{H}_{X,Y}^1(I,G)$.
2. Let $S \subseteq \text{an}_G(X \cup Y \cup I)$. Then $S \in \mathcal{A}_{X,Y}(I,R,G)$ if and only if S is an X,Y -separator in $\mathcal{H}_{X,Y}^1(I,G)$.
3. $S \in \mathcal{A}_{X,Y}^{\text{min}}(I,R,G)$ iff $S \in \mathcal{S}_{X,Y}(\mathcal{H}_{X,Y}^1(I,G))$.

In Theorem 4 we characterized separators between vertex sets as separators between singleton vertices in a modified undirected graph. Let $x \in X, y \in Y$. Let $\mathcal{H} \stackrel{\text{def}}{=} \mathcal{H}_{x,y}^1(I,G)$ be the graph that results from $\mathcal{H}_{x,y}^1(I,G)$ by adding all edges between x and $N_{\mathcal{H}_{x,y}^1(I,G)}(x)$, and all edges between y and $N_{\mathcal{H}_{x,y}^1(I,G)}(y)$. By Theorem 4:

$$\mathcal{A}_{x,y}^{\text{min}}(I,R,G) \stackrel{\text{Thm. 7}}{=} \mathcal{S}_{X,Y}(\mathcal{H}_{x,y}^1(I,G)) \stackrel{\text{Thm. 4}}{=} \mathcal{S}_{x,y}(\mathcal{H}) \quad (6)$$

In addition, if $S \subseteq \text{an}_G(I \cup X \cup Y)$ then:

$$\begin{aligned} S \in \mathcal{A}_{x,y}(I,R,G) &\Leftrightarrow S \text{ an } X,Y\text{-separator in } \mathcal{H}_{x,y}^1(I,G) \\ &\Leftrightarrow S \text{ an } x,y\text{-separator in } \mathcal{H} \end{aligned} \quad (7)$$

Eq. (6) and (7) allow us to reduce the problem of finding adjustment sets for X,Y in the DAG G to that of finding x,y -separators in the undirected graph \mathcal{H} .

4 ORDERING ADJUSTMENT SETS BY EFFICIENCY

Let $\mathcal{B}=(G,P)$ be a causal BN with observable variables $R \subseteq \text{V}(G)$, and let $X,Y,I \subseteq R$. We aim to estimate the interventional mean $\mathbb{E}_P(Y|\text{do}(\mathbf{X}), \mathbf{I})$. Using the method of covariate adjustment, the non-parametric estimator for the interventional mean depends on the adjustment set Z . Following previous work [Smucler et al., 2021, Smucler and

Rotnitzky, 2022], we consider unbiased estimators that converge, in distribution, to a normal distribution. We denote by $\sigma_Z^2(P)$ the variance of the normally-distributed estimator when computed using the adjustment set Z (see (4)). Different adjustment sets Z may yield estimators with varying levels of variance, making the choice of Z crucial for obtaining reliable and accurate estimates.

Let $Z_1, Z_2 \in \mathcal{A}_{X,Y}(I,R,G)$ be two valid adjustment sets for estimating the causal effect of X on Y in G . We say that Z_1 is *more efficient* than Z_2 , in notation $Z_1 \leq_G^\sigma Z_2$ if and only if $\sigma_{Z_1}^2(P) \leq \sigma_{Z_2}^2(P)$ for every joint probability distribution P that factorizes according to G . It has been established that \leq_G^σ does not induce a total order over the set of valid adjustment sets $\mathcal{A}_{X,Y}(I,R,G)$. In other words, there exist two distinct adjustment sets $Z_1, Z_2 \in \mathcal{A}_{X,Y}(I,R,G)$, such that $\sigma_{Z_1}^2(P) < \sigma_{Z_2}^2(P)$, and $\sigma_{Z_1}^2(P') > \sigma_{Z_2}^2(P')$, where P, P' are two distinct joint probability distributions that factorize according to G [Rotnitzky and Smucler, 2020]. In other words, there may be pairs of X,Y -adjustment sets in G that are incomparable with respect to efficiency.

Definition 5. Let $Z_1, Z_2 \in \mathcal{A}_{X,Y}(I,R,G)$. We say that Z_1 *dominates* Z_2 if $Z_1 \leq_G^\sigma Z_2$ and $|Z_1| \leq |Z_2|$, and that Z_1 *strictly dominates* Z_2 , if one of these inequalities is strict.

The Pareto-Optimal frontier of adjustment sets is the subset of sets in $\mathcal{A}_{X,Y}(I,R,G)$ that are not dominated.

Let H be an undirected graph, $A, B \subseteq \text{V}(H)$, and S_1, S_2 two (not necessarily minimal) A, B -separators in H . We denote by $S_1 \triangleleft_H S_2$ that S_1 separates A from $S_2 \setminus S_1$, and S_2 separates B from $S_1 \setminus S_2$ in H .

Theorem 8. [Smucler et al., 2021] Let $Z_1, Z_2 \in \mathcal{A}_{X,Y}(I,R,G)$, such that $Z_1, Z_2 \subseteq \text{V}(\mathcal{H}_{X,Y}^1(I,G))$. If $Z_1 \triangleleft_{\mathcal{H}_{X,Y}^1(I,G)} Z_2$, then $Z_1 \leq_G^\sigma Z_2$.

In Section A of the Appendix, we prove a result translating Theorem 8's efficiency criteria into a crucial structural property of separators for ranked enumeration.

Proposition 2. Let S_1, S_2 be two s, t -separators in H . Then $S_1 \triangleleft_H S_2$ if and only if $C_s(H-S_1) \subseteq C_s(H-S_2)$ and $C_t(H-S_2) \subseteq C_t(H-S_1)$.

4.1 PROOF OF THEOREM 1

We show how Proposition 1, Proposition 2, Theorem 6, and Theorem 8 are combined to prove Theorem 1.

Let $x \in X, y \in Y$, and \mathcal{H} the graph that results from $\mathcal{H}_{x,y}^1(I,G)$ by adding all edges between x and $N_{\mathcal{H}_{x,y}^1(I,G)}(x)$, and between y and $N_{\mathcal{H}_{x,y}^1(I,G)}(y)$. By Theorem 4, we have that $\mathcal{S}_{X,Y}(\mathcal{H}_{x,y}^1(I,G)) = \mathcal{S}_{x,y}(\mathcal{H})$. Let S_1, S_2 be X,Y -separators in $\mathcal{H}_{x,y}^1(I,G)$. From (7), we

have that $S_1 \preceq_{\mathcal{H}_{x,y}^1(I,G)} S_2$ if and only if $S_1 \preceq_{\mathcal{H}} S_2$. If, in addition, we have that S_1, S_2 are minimal X, Y -separators in $\mathcal{H}_{x,y}^1(I, G)$, then by Proposition 1, we have that $C_y(\mathcal{H}-S_1) \subseteq C_y(\mathcal{H}-S_2)$ if and only if $C_x(\mathcal{H}-S_2) \subseteq C_x(\mathcal{H}-S_1)$. By Proposition 2, we have that $S_1 \preceq_{\mathcal{H}} S_2$, and by Theorem 8 that

$$\text{If } C_y(\mathcal{H}-S_1) \subseteq C_y(\mathcal{H}-S_2), \text{ then } S_1 \leq_{\sigma}^G S_2. \quad (8)$$

Essentially, (8) translates the quality of adjustment sets in causal BNs with unobserved variables to a simple graphical property of minimal x, y -separators in an undirected graph. Therefore, the set of Pareto-Optimal adjustment sets in $\mathcal{A}_{x,y}^{\min}(I, R, G)$ precisely correspond to the important minimal x, y -separators in \mathcal{H} (Definition 2). By Theorem 6, the set of Pareto-Optimal adjustment sets whose size is at most k can be listed in time $O(k(n+m)4^k)$.

Proposition 2, and its consequence eq. (8) also generalize previous results of Smucler et al. [2021]. By Lemma 2, there exists a unique minimal x, y -separator S that is closest to Y . That is, $C_y(\mathcal{H}-S) \subseteq C_y(\mathcal{H}-S')$ for every $S' \in \mathcal{S}_{x,y}(\mathcal{H})$. From (8), this immediately translates to a unique, minimal, optimal, valid X, Y -adjustment set that can be found in polynomial time, thereby restoring the result of Smucler et al. [2021]. Theorem 5 established that there exists a unique, minimum x, y -separator S that is closest to y . From (8), this immediately translates to a unique, minimum, optimal, valid X, Y -adjustment set in the causal BN, that can be found in polynomial time, thereby restoring the result of Smucler et al. [2021].

5 RANKED ENUMERATION OF ALL s, t -SEPARATORS

We present, in Figure 2, an algorithm that lists all (not necessarily minimal) s, t -separators of G in ranked order by weight, where the secondary ranking is by distance to s . Formally, if S_1, S_2 are s, t -separators of G , then:

$$w(S_1) < w(S_2) \text{ or } (w(S_1) = w(S_2) \text{ and } S_1 \preceq S_2) \Rightarrow \quad (9) \\ S_1 \text{ is listed before } S_2$$

Due to space restrictions, the proofs of this section are deferred to Section C of the Appendix.

We begin by characterizing minimal s, t -separators that exclude a subset $U \subseteq \mathcal{V}(G)$ of vertices. We define: $\mathcal{S}_{s,t}(G, \bar{U}) \stackrel{\text{def}}{=} \{S \in \mathcal{S}_{s,t}(G) : S \subseteq \mathcal{V}(G) \setminus U\}$.

Theorem 9. $\mathcal{S}_{s,t}(G, \bar{U}) = \mathcal{S}_{s,t}(\text{Sat}(G, U))$.

Algorithm RankedEnumSeps of Figure 2 applies the Lawler technique with inclusion and exclusion constraints (e.g., see Golenberg et al. [2011]), leading to a simple polynomial-delay algorithm. By applying Theorem 5, it

easily follows that the delay of the algorithm is $O(Kn \cdot T(n, m))$ where K is the size of the largest s, t -separator printed by the algorithm, and $T(n, m)$ is the time to compute a minimum-weight s, t -separator.

Algorithm RankedEnumSeps($G, \{s, t\}$)

```

1: Compute  $S \in \mathcal{L}_{s,t}(G)$  closest to  $s$  {Thm. 5}
2:  $Q \leftarrow \text{PriorityQueue}(\text{card}, \preceq)$ 
3:  $Q.\text{push}(\langle G, S, I = \emptyset \rangle)$ 
4: while  $Q \neq \emptyset$  do
5:    $\langle G, S, I \rangle \leftarrow Q.\text{pop}()$ 
6:   print  $S$ 
7:   for all  $v_i \in S \setminus I = \{v_1, \dots, v_q\}$  do
8:      $I_i \leftarrow I \cup \{v_1, \dots, v_{i-1}\}$ 
9:      $H \leftarrow \text{Sat}(G, v_i)$  {Exclude  $v_i$ }
10:     $T \in \mathcal{L}_{s,t}(H - I_i)$  closest to  $s$  {Theorem 5}
11:     $Q.\text{push}(\langle H, T \cup I_i, I_i \rangle)$ 

```

Figure 2: Algorithm for enumerating all s, t -separators in ranked order by weight, and distance from s .

Theorem 10. Let S be an s, t -separator of G . There exists an s, t -separator S' printed by the algorithm where $S' \subseteq S$.

An immediate consequence of Theorem 10 is that every s, t -separator S is either printed or has a printed subset $S' \subset S$, ensuring no separator is lost. Our approach efficiently produces the K lightest s, t -separators first, making it a top- K enumeration algorithm for s, t -separators.

6 CONCLUSION

Previous methods optimized either adjustment set accuracy or size, but not both. Our approach integrates these aspects into an algorithmic framework for enumerating covariate adjustment sets that balances estimator variance and cost. We present three complementary algorithms suited to different practical and computational scenarios. The first efficiently enumerates all minimal adjustment sets of cost at most k that are optimal with respect to both cost and estimator variance. The second provides a polynomial-delay procedure for listing all valid adjustment sets, minimal and non-minimal alike, ranked by cost and variance, making it well-suited for broader exploration when adjustment sets may be large or completeness is required. The third algorithm bridges these two settings by enumerating all minimal adjustment sets of cost at most k in order of estimator quality, enabling robust selection in cases where optimal sets are sparse or cost functions are imprecise. Together, these results offer a flexible and principled approach to adjustment set selection under real-world constraints.

Acknowledgements

This work was supported by the Israel Science Foundation (ISF) Grant No. 1442/24.

References

- Anne Berry, Jean Paul Bordat, and Olivier Cogis. Generating all the minimal separators of a graph. *Int. J. Found. Comput. Sci.*, 11(3):397–403, 2000. doi: 10.1142/S0129054100000211. URL <https://doi.org/10.1142/S0129054100000211>.
- Daniel Bienstock. On the complexity of testing for odd holes and induced odd paths. *Discret. Math.*, 90(1):85–92, 1991. doi: 10.1016/0012-365X(91)90098-M. URL [https://doi.org/10.1016/0012-365X\(91\)90098-M](https://doi.org/10.1016/0012-365X(91)90098-M).
- Li Chen, Rasmus Kyng, Yang P. Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum flow and minimum-cost flow in almost-linear time. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 612–623, 2022. doi: 10.1109/FOCS54457.2022.00064.
- Nadia Creignou, Arne Meier, Julian-Steffen Müller, Johannes Schmidt, and Heribert Vollmer. Paradigms for parameterized enumeration. *Theory Comput. Syst.*, 60(4):737–758, 2017. doi: 10.1007/S00224-016-9702-4. URL <https://doi.org/10.1007/s00224-016-9702-4>.
- Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshantov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.
- Nicolas Derhy and Christophe Picouleau. Finding induced trees. *Discret. Appl. Math.*, 157(17):3552–3557, 2009. doi: 10.1016/j.dam.2009.02.009. URL <https://doi.org/10.1016/j.dam.2009.02.009>.
- Shimon Even and Guy Even. *Graph Algorithms, Second Edition*. Cambridge University Press, 2012.
- Shimon Even and Robert Endre Tarjan. Network flow and testing graph connectivity. *SIAM J. Comput.*, 4(4):507–518, 1975. doi: 10.1137/0204043. URL <https://doi.org/10.1137/0204043>.
- D. R. Ford and D. R. Fulkerson. *Flows in Networks*. Princeton University Press, USA, 2010. ISBN 0691146675.
- Konstantin Golenberg, Benny Kimelfeld, and Yehoshua Saviv. Optimizing and parallelizing ranked enumeration. *Proc. VLDB Endow.*, 4(11):1028–1039, 2011.
- Robert Haas and Michael Hoffmann. Chordless paths through three vertices. *Theor. Comput. Sci.*, 351(3):360–371, 2006. doi: 10.1016/j.tcs.2005.10.021. URL <https://doi.org/10.1016/j.tcs.2005.10.021>.
- Monika Rauch Henzinger, Satish Rao, and Harold N. Gabow. Computing vertex connectivity: New bounds from old techniques. *J. Algorithms*, 34(2):222–250, 2000.
- Arkady Kanevsky. On the number of minimum size separating vertex sets in a graph and how to find all of them. In *SODA*, pages 411–421. SIAM, 1990.
- Ton Kloks and Dieter Kratsch. Listing all minimal separators of a graph. *SIAM J. Comput.*, 27(3):605–613, 1998. doi: 10.1137/S009753979427087X. URL <https://doi.org/10.1137/S009753979427087X>.
- Kai-Yuan Lai, Hsueh-I Lu, and Mikkel Thorup. Three-in-a-tree in near linear time. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, pages 1279–1292, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450369794. doi: 10.1145/3357713.3384235. URL <https://doi.org/10.1145/3357713.3384235>.
- Eugene L. Lawler, Jan Karel Lenstra, and A. H. G. Rinnooy Kan. Generating all maximal independent sets: N_p -hardness and polynomial-time algorithms. *SIAM J. Comput.*, 9(3):558–565, 1980. doi: 10.1137/0209042. URL <https://doi.org/10.1137/0209042>.
- Dániel Marx. Important separators and parameterized algorithms. In Petr Kolman and Jan Kratochvíl, editors, *Graph-Theoretic Concepts in Computer Science - 37th International Workshop, WG 2011, Teplá Monastery, Czech Republic, June 21-24, 2011. Revised Papers*, volume 6986 of *Lecture Notes in Computer Science*, pages 5–10. Springer, 2011. doi: 10.1007/978-3-642-25870-1_2. URL https://doi.org/10.1007/978-3-642-25870-1_2.
- Judea Pearl. [bayesian analysis in expert systems]: Comment: Graphical models, causality and intervention. *Statistical Science*, 8(3):266–269, 1993. ISSN 08834237. URL <http://www.jstor.org/stable/2245965>.
- Judea Pearl. *Causality*. Cambridge University Press, Cambridge, UK, 2 edition, 2009. ISBN 978-0-521-89560-6. doi: 10.1017/CBO9780511803161.
- Emilija Perkovic, Johannes Textor, Markus Kalisch, and Marloes H. Maathuis. Complete graphical characterization and construction of adjustment sets in markov equivalence classes of ancestral graphs. *J. Mach. Learn. Res.*, 18:220:1–220:62, 2017. URL <http://jmlr.org/papers/v18/16-319.html>.

- Andrea Rotnitzky. Optimal adjustment sets in non-parametric graphical causal models. Talk presented at the Online Causal Inference Seminar, apr 2021. URL <https://drive.google.com/file/d/1WDFsU2dyaVrodyYi91nljBiwQogd1uAZ/view>. Slide 27.
- Andrea Rotnitzky and Ezequiel Smucler. Efficient adjustment sets for population average causal treatment effect estimation in graphical models. *J. Mach. Learn. Res.*, 21:188:1–188:86, 2020. URL <http://jmlr.org/papers/v21/19-1026.html>.
- Jakob Runge. Necessary and sufficient graphical conditions for optimal adjustment sets in causal graphical models with hidden variables. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 15762–15773, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/8485ae387a981d783f8764e508151cd9-Abstract.html>.
- Hong Shen and Weifa Liang. Efficient enumeration of all minimal separators in a graph. *Theoretical Computer Science*, 180(1):169–180, 1997. ISSN 0304-3975. doi: [https://doi.org/10.1016/S0304-3975\(97\)83809-1](https://doi.org/10.1016/S0304-3975(97)83809-1). URL <https://www.sciencedirect.com/science/article/pii/S0304397597838091>.
- E Smucler, F Sapienza, and A Rotnitzky. Efficient adjustment sets in causal graphical models with hidden variables. *Biometrika*, 109(1):49–65, 03 2021. ISSN 1464-3510. doi: [10.1093/biomet/asab018](https://doi.org/10.1093/biomet/asab018). URL <https://doi.org/10.1093/biomet/asab018>.
- Ezequiel Smucler and Andrea Rotnitzky. A note on efficient minimum cost adjustment sets in causal graphical models. *CoRR*, abs/2201.02037, 2022. URL <https://arxiv.org/abs/2201.02037>.
- Sara Mohammad Taheri, Vartika Tewari, Rohan Kapre, Ehsan Rahiminasab, Karen Sachs, Charles Tapley Hoyt, Jeremy Zucker, and Olga Vitek. Optimal adjustment sets for causal query estimation in partially observed biomolecular networks. *Bioinform.*, 39(Supplement-1):494–503, 2023. doi: [10.1093/BIOINFORMATICS/BTAD270](https://doi.org/10.1093/bioinformatics/btad270). URL <https://doi.org/10.1093/bioinformatics/btad270>.
- Ken Takata. Space-optimal, backtracking algorithms to list the minimal vertex separators of a graph. *Discret. Appl. Math.*, 158(15):1660–1667, 2010. doi: [10.1016/j.dam.2010.05.013](https://doi.org/10.1016/j.dam.2010.05.013). URL <https://doi.org/10.1016/j.dam.2010.05.013>.
- Benito van der Zander, Maciej Liskiewicz, and Johannes Textor. Constructing separators and adjustment sets in ancestral graphs. In Nevin L. Zhang and Jin Tian, editors, *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence, UAI 2014, Quebec City, Quebec, Canada, July 23-27, 2014*, pages 907–916. AUAI Press, 2014. URL https://dslpitt.org/uai/displayArticleDetails.jsp?mmnu=1&smnu=2&article_id=2527&proceeding_id=30.
- Benito van der Zander, Maciej Liskiewicz, and Johannes Textor. Separators and adjustment sets in causal graphs: Complete criteria and an algorithmic framework. *Artif. Intell.*, 270:1–40, 2019. doi: [10.1016/J.ARTINT.2018.12.006](https://doi.org/10.1016/j.artint.2018.12.006). URL <https://doi.org/10.1016/j.artint.2018.12.006>.

APPENDIX

A PROPERTIES OF SEPARATORS

PROPOSITION 1. *Let $S_1, S_2 \in \mathcal{S}_{s,t}(G)$. Then $C_s(G-S_1) \subseteq C_s(G-S_2)$ if and only if $C_t(G-S_2) \subseteq C_t(G-S_1)$.*

Proof. If $C_s(G-S_1) \subseteq C_s(G-S_2)$, then $C_s(G-S_1) \cup N_G(C_s(G-S_1)) \subseteq C_s(G-S_2) \cup N_G(C_s(G-S_2))$. By Lemma 1, we have that $S_1 = N_G(C_s(G-S_1))$. Therefore, $C_s(G-S_1) \cup S_1 \subseteq C_s(G-S_2) \cup N_G(C_s(G-S_2))$. In particular, $S_1 \cap C_t(G-S_2) = \emptyset$. This means that $C_t(G-S_2)$ is contained in the connected component of $G-S_1$ that contains t . By definition, $C_t(G-S_2) \subseteq C_t(G-S_1)$. The other direction is symmetrical. \square

PROPOSITION 2. *Let S_1, S_2 be two s, t -separators in H . Then $S_1 \triangleleft_H S_2$ if and only if $C_s(H-S_1) \subseteq C_s(H-S_2)$ and $C_t(H-S_2) \subseteq C_t(H-S_1)$.*

Proof. If $C_s(H-S_1) \subseteq C_s(H-S_2)$, then $\emptyset = C_s(H-S_2) \cap S_2 \supseteq C_s(H-S_1) \cap S_2$, and hence $C_s(H-S_1) \cap S_2 = \emptyset$. Consequently, $(S_2 \setminus S_1) \cap (C_s(H-S_1) \cup S_1) = \emptyset$. Every vertex connected to s in $H-S_1$ belongs to $C_s(H-S_1)$. Since $(S_2 \setminus S_1) \cap (C_s(H-S_1) \cup S_1) = \emptyset$, then S_1 separates s from $S_2 \setminus S_1$. Symmetrically, if $C_t(H-S_2) \subseteq C_t(H-S_1)$, then $(S_1 \setminus S_2) \cap (C_t(H-S_2) \cup S_2) = \emptyset$, thus S_2 separates t from $S_1 \setminus S_2$.

If S_1 separates s from $S_2 \setminus S_1$, then $(S_2 \setminus S_1) \cap C_s(H-S_1) = \emptyset$. By definition, $S_1 \cap C_s(H-S_1) = \emptyset$, and hence $S_2 \cap C_s(H-S_1) = \emptyset$. This, in turn, means that $C_s(H-S_1)$ is contained in the connected component of $H-S_2$ that contains s . By definition, $C_s(H-S_1) \subseteq C_s(H-S_2)$. Symmetrically, if S_2 separates t from $S_1 \setminus S_2$, then $C_t(H-S_2) \subseteq C_t(H-S_1)$. So, if $S_1 \triangleleft_H S_2$ then $C_s(H-S_1) \subseteq C_s(H-S_2)$ and $C_t(H-S_2) \subseteq C_t(H-S_1)$. \square

Proposition 3. Let $S \in \mathcal{S}_{s,t}(G)$ where $S \subseteq N_G(s)$. For every $T \in \mathcal{S}_{s,t}(G)$, it holds that $C_s(G-S) \subseteq C_s(G-T)$.

Proof. Since $S \subseteq N_G(s) \subseteq T \cup C_s(G-T)$, then $C_s(G-S) \subseteq C_s(G-T)$. \square

A.1 SEPARATORS BETWEEN VERTEX-SETS

In this Section, we prove Theorem 4 that follows from a series of Lemmas.

Lemma 3. Let A and B be two disjoint, non-adjacent subsets of $\mathcal{V}(G)$. Then $S \in \mathcal{S}_{A,B}(G)$ if and only if S is an A, B -separator, and for every $w \in S$, there exist two connected components $C_A, C_B \in \mathcal{C}(G-S)$ such that $C_A \cap A \neq \emptyset$, $C_B \cap B \neq \emptyset$, and $w \in N_G(C_A) \cap N_G(C_B)$.

Proof. If $S \in \mathcal{S}_{A,B}(G)$, then for every $w \in S$ it holds that $S \setminus \{w\}$ no longer separates A from B . Hence, there is a path from some $a \in A$ to some $b \in B$ in $G-(S \setminus \{w\})$. Let C_a and C_b denote the connected components of $\mathcal{C}(G-S)$ containing $a \in A$ and $b \in B$, respectively. Since C_a and C_b are connected in $G-(S \setminus \{w\})$, then $w \in N_G(C_a) \cap N_G(C_b)$.

Suppose that for every $w \in S$, there exist two connected components $C_A, C_B \in \mathcal{C}_G(S)$ such that $C_A \cap A \neq \emptyset$, $C_B \cap B \neq \emptyset$, and $w \in N_G(C_A) \cap N_G(C_B)$. If $S \notin \mathcal{S}_{A,B}(G)$, then $S \setminus \{w\}$ separates A from B for some $w \in S$. Since w connects C_A to C_B in $G-(S \setminus \{w\})$, no such $w \in S$ exists, and thus $S \in \mathcal{S}_{A,B}(G)$. \square

Observe that Lemma 3 implies Lemma 1. By Lemma 3, it holds that $S \in \mathcal{S}_{s,t}(G)$ if and only if S is an s, t -separator and $S \subseteq N_G(C_s(G-S)) \cap N_G(C_t(G-S))$. By definition, $N_G(C_s(G-S)) \subseteq S$ and $N_G(C_t(G-S)) \subseteq S$, and hence $S = N_G(C_s(G-S)) \cap N_G(C_t(G-S))$, and $S = N_G(C_s(G-S)) = N_G(C_t(G-S))$.

Lemma 4. Let G and H be graphs where $\mathcal{V}(G) = \mathcal{V}(H)$ and $\mathcal{E}(G) \subseteq \mathcal{E}(H)$. Let $A, B \subseteq \mathcal{V}(G)$ disjoint and non-adjacent. Let $S \in \mathcal{S}_{A,B}(G)$. If S is an A, B -separator in H , then $S \in \mathcal{S}_{A,B}(H)$.

Proof. Since $S \in \mathcal{S}_{A,B}(G)$, then by Lemma 3, for every $u \in S$ there exist two distinct connected components $C_A^u, C_B^u \in \mathcal{C}(G-S)$ where $C_A^u \cap A \neq \emptyset$, $C_B^u \cap B \neq \emptyset$, and $u \in N_G(C_A^u) \cap N_G(C_B^u)$. Since $\mathcal{E}(H) \supseteq \mathcal{E}(G)$, and since S is an A, B -separator in H , then $H-S$ contains two distinct connected components D_A^u, D_B^u where $C_A^u \subseteq D_A^u$ and $C_B^u \subseteq D_B^u$. Therefore, $w \in N_H(D_A^u) \cap N_H(D_B^u)$. By Lemma 3, we have that $S \in \mathcal{S}_{A,B}(H)$. \square

Lemma 5. Let $u \in \mathcal{V}(G) \setminus sB$ such that $N_G(u) \subseteq N_G(s)$. Then $\mathcal{S}_{s,B}(G) = \mathcal{S}_{s,B}(G-u)$

Proof. Let $S \in \mathcal{S}_{s,B}(G)$. We first show that $u \notin S$. Suppose, by way of contradiction, that $u \in S$. By Lemma 3, there exist two distinct vertices $x, y \in N_G(u)$ such that $x \in C_s(G-S)$ and $y \in C_B(G-S)$, where $C_B(G-S) \cap B \neq \emptyset$. By the assumption of the lemma that $N_G(u) \subseteq N_G(s)$, then $y \in N_G(s)$. But then, S is not an s, B -separator of G ; a contradiction. Hence $u \notin S$ for any $S \in \mathcal{S}_{s,B}(G)$.

Let $T \in \mathcal{S}_{s,B}(G-u)$. We show that T is an s, B -separator of G . If it is not, then since every s, B -path of $G-u$ is also an s, B -path of G , then $T \cup \{u\} \in \mathcal{S}_{s,B}(G)$. But this contradicts the fact that $u \notin S$ for every $S \in \mathcal{S}_{s,B}(G)$. Hence, T is an s, B -separator of G . By Lemma 4, we have that $T \in \mathcal{S}_{s,B}(G)$. Hence, we have that $\mathcal{S}_{s,B}(G-u) \subseteq \mathcal{S}_{s,B}(G)$. For the other direction, let $T \in \mathcal{S}_{s,B}(G)$. Clearly T is an s, B -separator of $G-u$. If $T \notin \mathcal{S}_{s,B}(G-u)$, then there exist a $T' \subset T$ s.t. $T' \in \mathcal{S}_{s,B}(G-u)$. By the previous direction, we have that $T' \in \mathcal{S}_{s,B}(G-u) \subseteq \mathcal{S}_{s,B}(G)$. But then, $T' \in \mathcal{S}_{s,B}(G)$ contradicting the minimality of T . Hence, $\mathcal{S}_{s,B}(G-u) = \mathcal{S}_{s,B}(G)$. \square

Lemma 6. Let $A \subseteq \mathcal{V}(G) \setminus Bs$. Let H be the graph that results from G by (1) adding all edges between s and $N_G(A)$, and (2) removing the vertices A and their adjacent edges from H . Then $\mathcal{S}_{sA,B}(G) = \mathcal{S}_{s,B}(H)$.

Proof. Let $T \in \mathcal{S}_{sA,B}(G)$, and let C_1, \dots, C_k denote the connected components of $\mathcal{C}(G-T)$ containing vertices from sA . By definition, $B \cap C_i = \emptyset$ for all $i \in \{1, 2, \dots, k\}$. Assume wlog that $s \in C_1$. Let H' denote the graph that results from G by adding all edges between s and $N_G(A)$. By definition, the edges added to G to form H' are between C_1 and $C_1 \cdots C_k \cup T$. Therefore, T separates sA from B in H' . Since $\mathbb{E}(H') \supseteq \mathbb{E}(G)$, then by Lemma 4, if $T \in \mathcal{S}_{sA,B}(G)$ and T is an sA, B -separator in H' , then $T \in \mathcal{S}_{sA,B}(H')$. Therefore, we have that $\mathcal{S}_{sA,B}(G) \subseteq \mathcal{S}_{sA,B}(H')$.

We now claim that $\mathcal{S}_{sA,B}(H') = \mathcal{S}_{s,B}(H')$. Take $T \in \mathcal{S}_{s,B}(H')$. We claim that T is an sA, B -separator in H' . Suppose it is not, and let $C \in \mathcal{C}(H'-T)$ such that $a, b \in C$ where $a \in A$ and $b \in B$. Let $y \in N_{H'}(a) \cap C$. By construction, $y \in N_{H'}(s)$. But then, $s \in N_{H'}(C)$ and hence T is not an s, B -separator in H' ; a contradiction. Since $T \in \mathcal{S}_{s,B}(H')$, then by Lemma 3, we have that for every $u \in T$ there exists a connected component $C_B^u \in \mathcal{C}(H'-T)$ s.t. $B \cap C_B^u \neq \emptyset$ and $u \in N_{H'}(C_s(H'-T)) \cap N_{H'}(C_B^u)$. By Lemma 3, we have that $T \in \mathcal{S}_{sA,B}(H')$. Hence $\mathcal{S}_{s,B}(H') \subseteq \mathcal{S}_{sA,B}(H')$. For the other direction, let $T \in \mathcal{S}_{sA,B}(H')$. Clearly, T is an s, B -separator of H' . If $T \notin \mathcal{S}_{s,B}(H')$, then there exists a $T' \subset T$ such that $T' \in \mathcal{S}_{s,B}(H')$. By the previous direction, we have that $T' \in \mathcal{S}_{sA,B}(H')$, but this contradicts the minimality of T . Hence, $\mathcal{S}_{s,B}(H') = \mathcal{S}_{sA,B}(H')$. Overall, we have shown that $\mathcal{S}_{sA,B}(G) \subseteq \mathcal{S}_{sA,B}(H') = \mathcal{S}_{s,B}(H')$.

Let $T \in \mathcal{S}_{s,B}(H')$. We first show that T separates sA from B in G ; if not, there is a path from $x \in sA$ to B in $G-T$. Let u be the first vertex on this path such that $u \notin sA$. Note that such a vertex $u \notin sA$ must exist because $B \cap sA = \emptyset$. In particular, $u \in N_G(sA)$, and by construction, $u \in N_{H'}(s)$. This means that there is a path from s to B (via u) in $H'-T$, which is a contradiction. Therefore, T is an sA, B -separator in G . If $T \notin \mathcal{S}_{sA,B}(G)$, then there is a $T' \in \mathcal{S}_{sA,B}(G)$ where $T' \subset T$. By the previous direction, $T' \in \mathcal{S}_{sA,B}(G) \subseteq \mathcal{S}_{s,B}(H')$, and hence $T' \in \mathcal{S}_{s,B}(H')$, contradicting the minimality of $T \in \mathcal{S}_{s,B}(H')$. Therefore, $\mathcal{S}_{sA,t}(G) = \mathcal{S}_{s,B}(H')$.

By construction, for every $u \in sA$, we have that $N_{H'}(u) \subseteq N_{H'}(s)$. From Lemma 5, we have that $\mathcal{S}_{s,B}(H') = \mathcal{S}_{s,B}(H)$. Therefore, $\mathcal{S}_{sA,t}(G) = \mathcal{S}_{s,B}(H)$. \square

THEOREM 4. Let $A, B \subseteq \mathcal{V}(G)$ be disjoint and non-adjacent, where $s \in A$ and $t \in B$. Let H be the graph that results from G by: (1) adding all edges between s and $N_G(A)$, (2) adding all edges between t and $N_G(B)$, and (3) removing vertices $AB \setminus \{s, t\}$ and their adjacent edges. Then $\mathcal{S}_{s,t}(H) = \mathcal{S}_{A,B}(G)$.

Proof. Let H_1 be the graph that results from G by adding all edges between s and $N_G(A)$, and removing vertices $A \setminus \{s\}$ from the graph. By Lemma 6, we have that $\mathcal{S}_{A,B}(G) = \mathcal{S}_{s,B}(H_1)$. By the assumption that A and B are disjoint and non-adjacent, then $N_G[B] = N_{H_1}[B]$. Now, let H_2 be the graph that results from H_1 by adding all edges between t and $N_{H_1}(B) = N_G(B)$, and removing vertices $B \setminus \{t\}$ from the graph H_2 . By Lemma 6, we have that $\mathcal{S}_{s,t}(H_2) = \mathcal{S}_{s,B}(H_1) = \mathcal{S}_{A,B}(G)$. \square

A.2 MINIMUM SEPARATORS

THEOREM 5. There exists a unique minimum s, t -separator $S^* \in \mathcal{L}_{s,t}(G)$ such that $S^* \preceq S$ for all $S \in \mathcal{L}_{s,t}(G)$, and S^* can be found in time $O(n \cdot T(n, m))$.

Theorem 5 is a straightforward extension of the following Theorem.

Theorem 11. (Cygan et al. [2015]) Let G be a non-weighted graph (i.e., $w(v) = 1$ for every $v \in V(G)$). There exists a unique minimum-cardinality s, t -separator $S^* \in \mathcal{L}_{s,t}(G)$ such that $S^* \preceq S$ for all $S \in \mathcal{L}_{s,t}(G)$, and S^* can be found in time $O(n \cdot T(n, m))$.

For completeness, we provide the proof of Theorem 5 herein.

Theorem 12. (Theorem 8.3 in Cygan et al. [2015]) For $X, Y \subseteq V(G)$. It holds that:

$$|N_G(X)| + |N_G(Y)| \geq |N_G(X \cap Y)| + |N_G(X \cup Y)|.$$

Proof Overview. The proof establishes that for every vertex $v \in V(G)$, the number of times it is accounted for in the left-hand-side (LHS) is at least as large as the number of times it is accounted for in the right-hand-side (RHS), thereby proving the claim. \square

Lemma 7. Let G be an undirected, weighted graph, with weight function $w : V(G) \rightarrow \mathbb{N}_{\geq 1}$. For $X, Y \subseteq V(G)$. It holds that:

$$w(N_G(X)) + w(N_G(Y)) \geq w(N_G(X \cap Y)) + w(N_G(X \cup Y)).$$

Proof Overview. The proof is identical to that of Theorem 12, establishing that for every vertex $v \in V(G)$, the number of times it is accounted for in the left-hand-side (LHS) is at least as large as the number of times it is accounted for in the right-hand-side (RHS), thereby proving the claim. Since the weights are positive, the claim follows. \square

Recall from Definition 1 that for two minimal s, t -separators $S_1, S_2 \in \mathcal{S}_{s,t}(G)$, it holds:

$$S_1 \preceq S_2 \text{ if and only if } C_s(G-S_1) \subseteq C_s(G-S_2).$$

Theorem 13. (Theorem 8.4 in Cygan et al. [2015]) Let G be an undirected, unweighted graph. There exists a minimum-cardinality s, t -separator $S^* \in \mathcal{L}_{s,t}(G)$, such that $S^* \preceq S$ for every $S \in \mathcal{L}_{s,t}(G)$.

Lemma 8 presents the weighted version of Theorem 13. The proof is similar to that of Theorem 13, and is provided below for completeness.

Lemma 8. Let G be an undirected, weighted graph, with weight function $w : V(G) \rightarrow \mathbb{N}_{\geq 1}$. There exists a minimum-weight s, t -separator $S^* \in \mathcal{L}_{s,t}(G)$, such that $S^* \preceq S$ for every $S \in \mathcal{L}_{s,t}(G)$.

Proof. Let $S_1, S_2 \in \mathcal{L}_{s,t}(G)$. By Lemma 7, and Lemma 1, we have that:

$$\begin{aligned} w(S_1) + w(S_2) &\stackrel{\text{Lemma 1}}{=} w(N(C_s(G-S_1))) + w(N(C_s(G-S_2))) \\ &\stackrel{\text{Lemma 7}}{\geq} w(N(C_s(G-S_1) \cap C_s(G-S_2))) + w(N(C_s(G-S_1) \cup C_s(G-S_2))). \end{aligned} \quad (10)$$

Define $S^- \stackrel{\text{def}}{=} N(C_s(G-S_1) \cap C_s(G-S_2))$ and $S^+ \stackrel{\text{def}}{=} N(C_s(G-S_1) \cup C_s(G-S_2))$. Since $s \in C_s(G-S_1) \cap C_s(G-S_2)$, and $t \notin C_s(G-S_1) \cup C_s(G-S_2)$, then both S^- and S^+ are s, t -separators of G . Therefore, $w(S^-) \geq \kappa_{s,t}(G) = w(S_1) = w(S_2)$, and $w(S^+) \geq \kappa_{s,t}(G) = w(S_1) = w(S_2)$.

From (10), we have that

$$2\kappa_{s,t}(G) = w(S_1) + w(S_2) \geq w(S^-) + w(S^+) \geq 2\kappa_{s,t}(G),$$

and hence, $w(S^-) = w(S^+) = \kappa_{s,t}(G)$. Since $S^- = N(C_s(G-S_1) \cap C_s(G-S_2))$, then by definition, $S^- \preceq S_1$ and $S^- \preceq S_2$. Since $\mathcal{L}_{s,t}(G)$, the set of minimum-weight s, t -separators of G , is finite, this proves the claim. \square

We are now ready to prove Theorem 5.

THEOREM 5. *There exists a unique minimum s, t -separator $S^* \in \mathcal{L}_{s,t}(G)$ such that $S^* \preceq S$ for all $S \in \mathcal{L}_{s,t}(G)$, and S^* can be found in time $O(n \cdot T(n, m))$.*

Proof. From Lemma 8, we have that $S^* \in \mathcal{L}_{s,t}(G)$ exists and is unique. We show that it can be found in time $O(n \cdot T(n, m))$. Finding a minimum-weight s, t -separator can be reduced, by standard techniques to the maximum-flow problem. Let $S_1 \in \mathcal{L}_{s,t}(G)$ be a minimum-weight s, t -separator found in this way. Now, we need to check whether there is another $S_2 \in \mathcal{L}_{s,t}(G)$ such that $S_2 \prec S_1$. If $C_s(G-S_2) \subset C_s(G-S_1)$, then by Proposition 1, it holds that $C_t(G-S_1) \subset C_t(G-S_2)$. In particular, $S_1 = N(C_t(G-S_1)) \subseteq C_t(G-S_2) \cup N(C_t(G-S_2)) = C_t(G-S_2) \cup S_2$. Since $S_1, S_2 \in \mathcal{L}_{s,t}(G)$, then $S_1 \not\subseteq S_2$, and hence $S_1 \cap C_t(G-S_2) \neq \emptyset$. In other words, if $S_2 \prec S_1$, then there must be a vertex $v \in S_1$ that belongs to $C_t(G-S_2)$. We check if this is the case by iterating over all vertices $v \in S_1$, and contracting $C_t(G-S_1) \cup \{v\}$ to the vertex t , and finding a minimum-weight s, t -separator in the resulting graph. If, for all $v \in S_1$, this results in a separator whose weight is strictly larger than $\kappa_{s,t}(G)$, then we have identified the minimum-weight s, t -separator that is closest to s . Otherwise, we repeat this procedure until no such vertex $v \in S_1$ is found – indicating that the computed s, t -separator is both minimum-weight, and closest to s . \square

B LISTING COST-CONSTRAINED MINIMAL s, t -SEPARATORS IN FPT-DELAY

In this section, we consider the case where the weight function $w : \mathcal{V}(G) \rightarrow \{1, \dots, c\}$ is bounded by a constant c . That is, $w(v) \leq c$ for all $v \in \mathcal{V}(G)$. We present an algorithm that given a threshold value W , returns all minimal s, t -separators $S \in \mathcal{S}_{s,t}(G)$ where $w(S) \leq W$ in ranked order by their distance from s . That is, if $S_1, S_2 \in \mathcal{S}_{s,t}(G)$ where $w(S_1) \leq W$, $w(S_2) \leq W$, and $S_1 \prec S_2$, then S_1 is printed before S_2 by the algorithm. To that end, we first transform G into an unweighted graph G' by using the common vertex-splitting technique that preserves G 's paths and connectivity structure. For every vertex $v \in \mathcal{V}(G)$, we generate a set C_v of $w(v) \leq c$ copies of v in G' and connect them to form a clique. Let $a', b' \in \mathcal{V}(G')$ be copies of distinct vertices $a, b \in \mathcal{V}(G)$, respectively. Then $(a', b') \in \mathbb{E}(G')$ if and only if $(a, b) \in \mathbb{E}(G)$. It is easily seen that $S \in \mathcal{S}_{s,t}(G)$ if and only if $S' \in \mathcal{S}_{s,t}(G')$, where $S' \stackrel{\text{def}}{=} \cup_{v \in S} C_v$. Note also that $|S'| = w(S)$. Hence, the problem is reduced to that of returning all minimal s, t -separators of G' whose *cardinality* is at most W in ranked order by their distance from s . The rest of this section is devoted to this problem.

For any natural number $k \geq 0$, we denote by $\mathcal{S}_{s,t,k}(G)$ the minimal s, t -separators whose cardinality is at most k .

$$\mathcal{S}_{s,t,k}(G) \stackrel{\text{def}}{=} \{S \in \mathcal{S}_{s,t}(G) : |S| \leq k\} \quad (11)$$

Let $C_1, C_2 \subseteq \mathcal{V}(G)$. We say that they are *incomparable* if $C_1 \not\subseteq C_2$ and $C_2 \not\subseteq C_1$. Recall Definition 1 of the partial order \preceq between minimal s, t -separators. In Figure 3, we present the enumeration algorithm that lists all minimal s, t -separators whose size is at most k (i.e., $\mathcal{S}_{s,t,k}(G)$) according to the order \prec . Recall that $\mathcal{S}_{s,t,k}^*(G)$ is the set of important minimal s, t -separators of size at most k (see Section 2.3).

B.1 OVERVIEW OF PROOF OF CORRECTNESS

Correctness is established by showing that a subset $S \subseteq \mathcal{V}(G)$ is printed by the algorithm if and only if $S \in \mathcal{S}_{s,t}(G)$, and $|S| \leq k$ (i.e., $S \in \mathcal{S}_{s,t,k}(G)$). This is established by Theorems 14 and 15. In Theorem 16, we show that if $S_1, S_2 \in \mathcal{S}_{s,t,k}(G)$, and $S_1 \prec S_2$, then S_1 is printed before S_2 . Finally, we establish FPT-delay in Theorem 17. The results follow from a series of lemmas.

Theorem 14. If $S \subseteq \mathcal{V}(G)$ is printed, then $S \in \mathcal{S}_{s,t,k}(G)$, and S is printed exactly once.

Theorem 15. Let $T \in \mathcal{S}_{s,t,k}(G)$. Then T is printed by SmallMinimalSeps in Figure 3.

Theorem 14 and Theorem 15 together establish that a subset $T \subseteq \mathcal{V}(G)$ is printed by the algorithm if and only if $T \in \mathcal{S}_{s,t,k}(G)$. In other words, SmallMinimalSeps prints precisely the set $\mathcal{S}_{s,t,k}(G)$.

Theorem 16. Let $S_1, S_2 \in \mathcal{S}_{s,t,k}(G)$. If $S_1 \prec S_2$, then S_1 is printed before S_2 by Algorithm SmallMinimalSeps.

Theorem 16 establishes that the minimal s, t -separators are printed in an order that corresponds to the desirability of the adjustment sets, in terms of their variance (see Section 4).

Theorem 17. The delay between the printing of minimal s, t -separators whose size is at most k is $O(k^2 4^k (n + m))$.

Lemma 9. Let $v \in N_G(s)$, and let G' denote the graph that results from G by contracting the edge (s, v) to s . Then $\mathcal{S}_{s,t}(G') = \{S \in \mathcal{S}_{s,t}(G) : v \notin S\}$.

Algorithm SmallMinimalSep(G, s, t, k)

Input: Graph G , and $s, t, \in \mathcal{V}(G)$.
Output: $\mathcal{S}_{s,t,k}(G)$.

- 1: **if** $(s, t) \in \mathbb{E}(G)$ **then**
- 2: Print \perp
- 3: **return**
- 4: $Q \leftarrow \text{PriorityQueue}(\preceq), \mathcal{M} \leftarrow \emptyset$
- 5: **Compute** $\mathcal{S}_{s,t,k}^*(G)$ {Theorem 6}
- 6: **for** $S \in \mathcal{S}_{s,t,k}^*(G)$ **do**
- 7: $Q.\text{push}(S)$
- 8: **while** Q is not empty **do**
- 9: $S \leftarrow Q.\text{pop}()$
- 10: Print S
- 11: $\mathcal{M}.\text{push}(S)$
- 12: Define $H_S: \mathcal{V}(H_S) = \{s\} \cup S \cup C_t(G-S),$
 $\mathbb{E}(H_S) = \mathbb{E}(G[\mathcal{V}(H_S)]) \cup \{(s, v) : v \in S\}$
- 13: **for** $v \in S$ **do**
- 14: Let H_S^v be the graph that results from H_S by contracting the
edge (s, v) to s .
- 15: **Compute** $\mathcal{S}_{s,t,k}^*(H_S^v)$ {Theorem 6}
- 16: **for** $T \in \mathcal{S}_{s,t,k}^*(H_S^v)$ **do**
- 17: **if** $T \notin Q$ AND $T \notin \mathcal{M}$ **then**
- 18: $Q.\text{push}(T)$

Figure 3: Algorithm for listing the minimal s, t -separators of G whose size is at most k , ranked by \preceq .

Proof. Let G'' be the graph that results from G by adding all edges between s and $N_G(v)$. By definition, this means that $N_{G''}(v) \subseteq N_{G''}(s)$. We first show that $\mathcal{S}_{s,t}(G'') = \{S \in \mathcal{S}_{s,t}(G) : v \notin S\}$.

Let $S \in \mathcal{S}_{s,t}(G)$ such that $v \notin S$. Since $v \in N_G(s)$, then $v \in C_s(G-S)$, and hence $N_G(v) \subseteq S \cup C_s(G-S)$. Therefore, S is an s, t -separator in G'' as well. Since $\mathbb{E}(G'') \supseteq \mathbb{E}(G)$, then by Lemma 4, $S \in \mathcal{S}_{s,t}(G'')$.

Now, let $T \in \mathcal{S}_{s,t}(G'')$. Since $\mathbb{E}(G) \subseteq \mathbb{E}(G'')$ then clearly T is an s, t -separator of G . Since $N_{G''}(v) \subseteq N_{G''}(s)$, then by Lemma 5, it holds that $\mathcal{S}_{s,t}(G'') = \mathcal{S}_{s,t}(G''-v)$. Therefore, we have that $v \notin T$. If $T \notin \mathcal{S}_{s,t}(G)$, then there exists a $T' \subset T$ such that $T' \in \mathcal{S}_{s,t}(G)$. Since $v \notin T$, then $v \notin T'$. We have previously established that $\mathcal{S}_{s,t}(G'') \supseteq \{S \in \mathcal{S}_{s,t}(G) : v \notin S\}$, and hence $T' \in \mathcal{S}_{s,t}(G'')$. But this contradicts the minimality of T . Therefore, $T \in \{S \in \mathcal{S}_{s,t}(G) : v \notin S\}$, and we have that $\mathcal{S}_{s,t}(G'') = \{S \in \mathcal{S}_{s,t}(G) : v \notin S\}$.

By construction, we have that $N_{G''}(v) \subseteq N_{G''}(s)$. By Lemma 5, we have that $\mathcal{S}_{s,t}(G'') = \mathcal{S}_{s,t}(G''-v) = \mathcal{S}_{s,t}(G')$. Therefore, we get that $\mathcal{S}_{s,t}(G') = \{S \in \mathcal{S}_{s,t}(G) : v \notin S\}$. \square

Lemma 10. Let $S, T \in \mathcal{S}_{s,t}(G)$. Then:

$$C_s(G-S) \subseteq C_s(G-T) \text{ if and only if } T \subseteq S \cup C_t(G-S).$$

Proof. If $T \subseteq S \cup C_t(G-S)$, then by definition $T \cap C_s(G-S) = \emptyset$. Therefore, $C_s(G-S)$ remains connected in $G-T$. This means that $C_s(G-S) \subseteq C_s(G-T)$.

Now, suppose that $C_s(G-S) \subseteq C_s(G-T)$. By Lemma 1, it holds that $S = N_G(C_s(G-S))$. Since $C_s(G-S) \subseteq C_s(G-T)$, then $S = N_G(C_s(G-S)) \subseteq T \cup C_s(G-T)$. Since $S \subseteq T \cup C_s(G-T)$ then by definition it holds that $S \cap C_t(G-T) = \emptyset$. This, in turn, implies that $C_t(G-T)$ remains connected in $G-S$. In particular, we have that $C_t(G-T) \subseteq C_t(G-S)$. By Lemma 1, it holds that $T = N_G(C_t(G-T))$. Since $C_t(G-T) \subseteq C_t(G-S)$, then $T = N_G(C_t(G-T)) \subseteq S \cup C_t(G-S)$. \square

Lemma 11. Let $S \in \mathcal{S}_{s,t}(G)$, and let H_S be the graph that results from G by adding all edges from s to S . That is, $\mathbb{E}(H_S) = \mathbb{E}(G) \cup \{(s, v) : v \in S\}$. Then:

$$\mathcal{S}_{s,t}(H_S) = \{Q \in \mathcal{S}_{s,t}(G) : Q \subseteq S \cup C_t(G-S)\}$$

Proof. Let $Q \in \mathcal{S}_{s,t}(G)$ where $Q \subseteq S \cup C_t(G-S)$. Since $Q \cap C_s(G-S) = \emptyset$, then $C_s(G-S)$ remains connected in $G-Q$. Therefore, $C_s(G-S) \subseteq C_s(G-Q)$. By Lemma 1, $S = N_G(C_s(G-S))$. Since $C_s(G-S) \subseteq C_s(G-Q)$, then $S = N_G(C_s(G-S)) \subseteq C_s(G-Q) \cup Q$. In particular, $S \cap C_t(G-Q) = \emptyset$. Consequently, Q separates $C_t(G-Q)$ from s in H_S as well. That is, Q is an s, t -separator in H_S . Since $\mathbb{E}(H_S) \supseteq \mathbb{E}(G)$, then $Q \in \mathcal{S}_{s,t}(H_S)$.

Let $T \in \mathcal{S}_{s,t}(H_S)$. By construction, $S \in \mathcal{S}_{s,t}(H_S)$ where $S \subseteq N_H(s)$. By Proposition 3, $C_s(H_S-S) \subseteq C_s(H_S-T)$. By Lemma 10, it holds that $T \subseteq S \cup C_t(H_S-S)$. Since, by construction, $C_t(H_S-S) = C_t(G-S)$, we get that $T \subseteq S \cup C_t(G-S)$. \square

Lemma 12. Let $T \in \mathcal{S}_{s,t,k}(G)$. Exactly one of the following holds: (1) $T \in \mathcal{S}_{s,t,k}^*(G)$ or (2) There exists a minimal s, t -separator $S \in \mathcal{S}_{s,t,k}^*(G)$ such that $S \prec T$.

Proof. By induction on $|C_s(G-T)|$. If $|C_s(G-T)| = 1$, then clearly $T \subseteq N_G(s)$. By Lemma 2, T is the unique minimal s, t -separator that is closest to s , and hence $T \in \mathcal{S}_{s,t,k}^*(G)$. So, we assume that the claim holds for all $T \in \mathcal{S}_{s,t,k}(G)$, where $1 \leq |C_s(G-S)| \leq \ell$. Let $T \in \mathcal{S}_{s,t,k}(G)$, where $|C_s(G-S)| = \ell + 1$. If $T \in \mathcal{S}_{s,t,k}^*(G)$, then we are done. Otherwise, if $T \notin \mathcal{S}_{s,t,k}^*(G)$, then since $|T| \leq k$, it must hold that $T \notin \mathcal{S}_{s,t}^*(G)$. By definition 2, there exists a $T' \in \mathcal{S}_{s,t}(G)$ such that $T' \prec T$ (i.e., $C_s(G-T') \subset C_s(G-T)$), and $|T'| \leq |T| \leq k$. Consequently, $|C_s(G-T')| < |C_s(G-T)| = \ell + 1$, and $|C_s(G-T')| \leq \ell$. Since $T' \in \mathcal{S}_{s,t,k}(G)$ and $|C_s(G-T')| \leq \ell$, then by the induction hypothesis, either $T' \in \mathcal{S}_{s,t,k}^*(G)$, in which case $T' \prec T$, thus proving the claim. Otherwise, there exists an $S \in \mathcal{S}_{s,t,k}^*(G)$ such that $S \prec T'$. Hence, $S \prec T' \prec T$, and $S \prec T$, thus proving the claim. \square

Lemma 13. Let $T \in \mathcal{S}_{s,t,k}(G)$. There exists a $S \in \mathcal{S}_{s,t,k}^*(G)$ such that $S \preceq T$, and $T \subseteq S \cup C_t(G-S)$.

Proof. If $T \in \mathcal{S}_{s,t,k}^*(G)$, then the claim is immediate. If $T \notin \mathcal{S}_{s,t,k}^*(G)$ then, by Lemma 12, there exists an $S \in \mathcal{S}_{s,t,k}^*(G)$, such that $S \prec T$. By Lemma 10, $T \subseteq S \cup C_t(G-S)$. \square

THEOREM 14. If $S \subseteq V(G)$ is printed, then $S \in \mathcal{S}_{s,t,k}(G)$, and S is printed exactly once.

Proof. Every subset of vertices inserted into the queue (in lines 7 and 18) is pushed exactly once and has cardinality at most k . Therefore, we only need to show that every subset of vertices pushed into the queue Q , and printed by the algorithm, belongs to $\mathcal{S}_{s,t}(G)$. Suppose, by way of contradiction, that this is not the case, and let $T \subseteq V(G)$ be the first subset of vertices printed where $T \notin \mathcal{S}_{s,t}(G)$. Then T must be inserted into the queue in line 18. Consider the set S that was printed before T is inserted into the queue. By our assumption $S \in \mathcal{S}_{s,t}(G)$. Therefore, $T \in \mathcal{S}_{s,t,k}^*(H_S^v)$, where $v \in S$. By Lemma 11, $\mathcal{S}_{s,t}(H_S) \subseteq \mathcal{S}_{s,t}(G)$. Since $v \in N_{H_S}(s)$, and H_S^v is the graph that results from H_S by contracting the edge (s, v) to vertex v , by Lemma 9, it holds that $\mathcal{S}_{s,t}(H_S^v) \subseteq \mathcal{S}_{s,t}(H_S) \subseteq \mathcal{S}_{s,t}(G)$. Since $T \in \mathcal{S}_{s,t,k}^*(H_S^v) \subseteq \mathcal{S}_{s,t}(H_S)$, we get that $T \in \mathcal{S}_{s,t}(G)$, which brings us to a contradiction. \square

THEOREM 15. Let $T \in \mathcal{S}_{s,t,k}(G)$. Then T is printed by SmallMinimalSeps in Figure 3.

Proof. If $T \in \mathcal{S}_{s,t,k}^*(G)$, then T is inserted into the queue in line 7, and will be printed. Therefore, assume that $T \notin \mathcal{S}_{s,t,k}^*(G)$. Suppose that T is not printed. Let $T' \in \mathcal{S}_{s,t}(G)$ be the largest minimal s, t -separator, with respect to \prec , that is printed by the algorithm, such that $T' \preceq T$. In other words, there does not exist a $T'' \in \mathcal{S}_{s,t}(G)$, that is printed by the algorithm where $T' \prec T'' \preceq T$. By Lemma 13, and the fact that $T \notin \mathcal{S}_{s,t,k}^*(G)$ such a separator T' exists.

Since $C_s(G-T') \subset C_s(G-T)$, then by Lemma 10, it holds that $T \in T' \cup C_t(G-T')$. By Lemma 11, it holds that $T \in \mathcal{S}_{s,t}(H_{T'})$. Consider what happens when T' is popped from the queue in line 9, and the graph $H_{T'}$ is generated in line 12. Since $T \neq T'$ (we assume that T is not printed), $T' \subseteq N_{H_{T'}}(s)$, and $T \in \mathcal{S}_{s,t}(H_{T'})$, then there exists a vertex $v \in T'$, such that $T \in \mathcal{S}_{s,t}(H_{T'}^v)$ (see line 14). If $T \in \mathcal{S}_{s,t,k}^*(H_{T'}^v)$, then T is pushed into the queue in line 18, and will therefore be printed. Otherwise, by Lemma 13, there exists an $S \in \mathcal{S}_{s,t,k}^*(H_{T'}^v)$, such that $C_s(H_{T'}^v-S) \subseteq C_s(H_{T'}^v-T)$. By construction, we have that $C_s(H_{T'}-T') \subset C_s(H_{T'}^v-S) \subseteq C_s(H_{T'}-T)$. Since S is pushed into the queue in line 18, then it

will be printed by the algorithm in line 10. By Theorem 14, we have that $S \in \mathcal{S}_{s,t,k}(G)$ is printed by the algorithm, where $T' \prec S \preceq T$, contradicting our assumption that T' is maximal with respect to the partial order \prec . \square

THEOREM 16. *Let $S_1, S_2 \in \mathcal{S}_{s,t,k}(G)$. If $S_1 \prec S_2$, then S_1 is printed before S_2 by Algorithm SmallMinimalSeps.*

Proof. By Theorem 15, both S_1 and S_2 are printed by the algorithm. Consider the point in time where S_2 is pushed into the queue Q .

1. **Case 1:** $S_1 \in \mathcal{M}$. In that case, when S_2 is pushed into the queue, S_1 has already been printed, and hence S_1 is printed before S_2 .
2. **Case 2:** $S_1 \in Q$. Since Q is a priority queue sorted according to \prec , then S_1 will be popped from the queue Q (in line 9), and printed (in line 10) before S_2 is popped (and printed).
3. **Case 3:** S_1 is generated and inserted into the queue *after* S_2 is printed. In that case, by the workings of the algorithm, $S_1 \in \mathcal{S}_{s,t,k}(H_{S_2}^v)$ for some $v \in S_2$ (see lines 13–18). By Lemma 9, $S_1 \in \mathcal{S}_{s,t,k}(H_{S_2}^v) \subseteq \mathcal{S}_{s,t,k}(H_{S_2})$. By Lemma 11, if $S_1 \in \mathcal{S}_{s,t,k}(H_{S_2})$, then $S_1 \in \mathcal{S}_{s,t,k}(G)$ where $S_1 \subseteq S_2 \cup C_t(G-S_2)$. By Lemma 10, we have that $C_s(G-S_2) \subseteq C_s(G-S_1)$; a contradiction. Therefore, only cases 1 and 2 are possible, which means that S_1 is printed before S_2 . \square

THEOREM 17. *The delay between the printing of minimal s, t -separators whose size is at most k is $O(k^2 4^k (n + m))$.*

Proof. The size of the queue Q and the data structure \mathcal{M} , can be at most n^k . We make the standard assumption that these data structures allow logarithmic insertion and extraction, which take time $O(k \log n)$. Applying Theorem 6, which states that there are at most 4^k important separators that can be found in time $O(k 4^k (n + m))$, we get that the loop in lines (13)-(18) runs in time: $O(k \cdot (n + 4^k \cdot k \cdot (n + m)) + k \cdot 4^k \cdot \log n)$. Overall, the delay is $O(4^k k^2 (n + m))$. \square

C PROOFS FROM SECTION 5

We prove that $\mathcal{S}_{s,t}(G, \bar{U}) = \mathcal{S}_{s,t}(\text{Sat}(G, U))$. We proceed by a series of lemmas.

Lemma 14. Let $u \in \mathcal{V}(G)$ such that $N_G[u]$ forms a clique. Then $u \notin S$ for every $S \in \mathcal{S}_{s,t}(G)$.

Proof. Let $S \in \mathcal{S}_{s,t}(G)$. By Lemma 1, $G-S$ contains two full connected components $C_s(G-S)$ and $C_t(G-S)$ containing s and t respectively, such that $S = N_G(C_s(G-S)) = N_G(C_t(G-S))$. Therefore, if $u \in S$, then it has two neighbors $v_1 \in C_s(G-S)$ and $v_2 \in C_t(G-S)$ that are connected by an edge (because $N_G[u]$ is a clique). But then, there is an s, t -path in $G-S$ that avoids S , which contradicts the fact that S is an s, t -separator. \square

Lemma 15. If $S \in \mathcal{S}_{s,t}(G, \bar{u})$, there exists a connected component $C_u \in \mathcal{C}(G-S)$ such that $N_G[u] \subseteq C_u \cup S$.

Proof. Let $C_u \in \mathcal{C}(G-S)$ be the connected component that contains u . Such a component must exist because, by Lemma 14, $u \notin S$. If $N_G(u) \not\subseteq C_u \cup S$, then there exists a vertex $v \in N_G(u)$ that resides in a connected component $C_v \in \mathcal{C}(G-S)$ distinct from C_u . But this is a contradiction because, by definition, $(u, v) \in E(G)$. Hence, $C_v = C_u$, and this proves the claim. \square

Lemma 16. Let $u \in \mathcal{V}(G)$. Then $\mathcal{S}_{s,t}(G, \bar{u}) = \mathcal{S}_{s,t}(\text{Sat}(G, \{u\}))$.

Proof. Let $S \in \mathcal{S}_{s,t}(G, \bar{u})$. By Lemma 15, there exists a connected component $C_u \in \mathcal{C}(G-S)$ that contains u , where $N_G[u] \subseteq C_u \cup S$. Therefore, no added edge in $E(\text{Sat}(G, \{u\})) \setminus E(G)$ connects vertices in distinct connected components in $\mathcal{C}(G-S)$. Hence, S separates s and t also in $\text{Sat}(G, \{u\})$. Since the addition of edges cannot eliminate any path between s and t , we get that S is a minimal s, t -separator also in $\text{Sat}(G, \{u\})$. Hence, $\mathcal{S}_{s,t}(G, \bar{u}) \subseteq \mathcal{S}_{s,t}(\text{Sat}(G, \{u\}))$.

Now, let $S \in \mathcal{S}_{s,t}(\text{Sat}(G, \{u\}))$. Hence, $N_G[u]$ is a clique in $\text{Sat}(G, \{u\})$. By Lemma 14, $u \notin S$. Since G is a subgraph of $\text{Sat}(G, \{u\})$, then if S separates s from t in $\text{Sat}(G, \{u\})$, it must separate s from t in G . Hence, S is an s, t -separator in G where $u \notin S$. It is left to show that S is a *minimal* s, t -separator in G . Suppose that it is not, and let $S' \subset S$ be

a minimal s, t -separator of G . Since $u \notin S$, then $u \notin S'$. By definition, $S' \in \mathcal{S}_{s,t}(G, \bar{u})$. By the previous direction, $S' \in \mathcal{S}_{s,t}(G, \bar{u}) \subseteq \mathcal{S}_{s,t}(\text{Sat}(G, \{u\}))$, and hence $S' \in \mathcal{S}_{s,t}(\text{Sat}(G, \{u\}))$. But this is a contradiction to the minimality of S . Therefore, $\mathcal{S}_{s,t}(\text{Sat}(G, \{u\})) \subseteq \mathcal{S}_{s,t}(G, \bar{u})$, and this completes the proof. \square

THEOREM 9. $\mathcal{S}_{s,t}(G, \bar{U}) = \mathcal{S}_{s,t}(\text{Sat}(G, U))$.

Proof. The fact that $\mathcal{S}_{s,t}(G, \bar{U}) = \mathcal{S}_{s,t}(\text{Sat}(G, U))$ follows from Lemma 16 by induction on $|U|$. \square

Let $0 \leq k \leq n$ be an integer, and $\mathcal{S}_{s,t}(G, \bar{U})^k$ and $\mathcal{S}_{s,t}(\text{Sat}(G, U))^k$ denote the sets of minimal s, t -separators in $\mathcal{S}_{s,t}(G, \bar{U})$ and $\mathcal{S}_{s,t}(\text{Sat}(G, U))$ whose size is exactly k , respectively. Since $\mathcal{S}_{s,t}(G, \bar{U}) = \mathcal{S}_{s,t}(\text{Sat}(G, U))$, then $\mathcal{S}_{s,t}(G, \bar{U})^k = \mathcal{S}_{s,t}(\text{Sat}(G, U))^k$ for every integer $0 \leq k \leq n$. In particular, this is the case for $k = \kappa_{s,t}(G, \bar{U}) = \kappa_{s,t}(\text{Sat}(G, U))$. Hence, $\mathcal{L}_{s,t}(G, \bar{U}) = \mathcal{L}_{s,t}(\text{Sat}(G, U))$.

THEOREM 10. *Let S be an s, t -separator of G . There exists an s, t -separator S' printed by the algorithm where $S' \subseteq S$.*

Proof. Let T be an s, t -separator of G , and suppose, by way of contradiction, that neither T , nor any of its subsets are printed. Every triple $\langle H, S, I \rangle$ pushed into the queue Q in lines 3 and 11 corresponds to a pair of inclusion/exclusion constraints that restrict the set of s, t -separators to those that include vertices I , and exclude vertices $U \subseteq V(G)$ that have been saturated in G (i.e., to form H). Let $\langle H, S, I \rangle$ be the triple, inserted into Q , where: (1) $I \subseteq T$, and (2) $U \subseteq V(G) \setminus T$, which maximizes $|I| + |U|$. Note that such a triple $\langle H, S, I \rangle$ must exist because the first triple pushed into the queue Q in line 3 is $\langle G, S, \emptyset \rangle$ where $S \in \mathcal{L}_{s,t}(G)$, $I = \emptyset \subseteq T$, and no vertex of G has yet been saturated and hence $U = \emptyset \subseteq V(G) \setminus T$.

Let $S \setminus I = \{v_1, \dots, v_q\}$. By our assumption, $S \not\subseteq T$. Let $\ell \leq q$ be the smallest index such that $v_\ell \notin T$. In other words, $\{v_1, \dots, v_{\ell-1}\} \subseteq T$, and $v_\ell \notin T$. In the ℓ th iteration of the loop in lines 7–11, the algorithm generates a triple $\langle H_\ell, S_\ell, I_\ell \rangle$, where $I_\ell \stackrel{\text{def}}{=} I \cup \{v_1, \dots, v_{\ell-1}\} \subseteq T$, and H_ℓ is the graph that, by Theorem 9, materializes the condition of excluding $U \cup \{v_\ell\}$. In other words, the algorithm generates a triple with inclusion constraints $I \subseteq I_\ell \subseteq T$, and exclusion constraint $U_\ell \stackrel{\text{def}}{=} U \cup \{v_\ell\} \supset U$, where $U_\ell \subseteq V(G) \setminus T$, and $|U_\ell| > |U|$. But then, $\langle H, S, I \rangle$ does not maximize $|I| + |U|$; a contradiction. \square

D MINIMAL SEPARATORS AND CHORDLESS s, t -PATHS

In this section we show that given a set $I \subseteq V(G)$, it is NP-hard to decide whether there exists a minimal s, t -separator $S \in \mathcal{S}_{s,t}(G)$ such that $I \subset S$. We prove this by showing a reduction from the problem 3-IN-A-PATH that asks whether there is an induced (or chordless) path containing three given terminals. Bienstock [1991] has shown that deciding whether two terminals belong to an induced cycle is NP-hard. From this, it is easy to show that the 3-IN-A-PATH problem is NP-hard even for graphs whose degree is at most three [Derhy and Picouleau, 2009]. In fact, even deciding whether there is such a path of length at most k was shown to be $W[1]$ -complete with respect to the length parameter k [Haas and Hoffmann, 2006]. The related problem, called THREE-IN-A-TREE, for deciding whether there is an induced tree containing three terminals, is in PTIME [Lai et al., 2020].

Theorem 18. *Let $v \in V(G)$. There exists a minimal s, t -separator that includes v if and only if there exists a chordless s, t -path through v .*

Proof. Let $S \in \mathcal{S}_{s,t}(G)$ where $v \in S$, and let $C_s(G-S)$, $C_t(G-S)$ denote the connected components of $G-S$ that contain s and t respectively. By Lemma 1, there exists a path from s to v where all the internal vertices belong to $C_s(G-S)$. Let P_{sv} denote the shortest such path. Likewise, let P_{vt} denote the shortest path from v to t where all internal vertices belong to $C_t(G-S)$. Clearly, P_{sv} and P_{vt} are both chordless paths. Since $C_s(G-S) \cap C_t(G-S) = \emptyset$, then $V(P_{sv}) \cap V(P_{vt}) = \{v\}$. Since $S \in \mathcal{S}_{s,t}(G)$, then there are no edges between vertices in $C_s(G-S)$ and vertices in $C_t(G-S)$. Consequently, there are no edges between vertices in $V(P_{sv})$ and $V(P_{vt})$. Therefore, the path $P_{sv}P_{vt}$ is a chordless s, t -path that passes through v . In other words, if $v \in S$, then there is an induced s, t -path through v .

Let $P = s, a_1, \dots, a_k, v, b_1, \dots, b_\ell, t$ denote a simple, chordless s, t -path through v . If $v \in N_G(s)$ ($v \in N_G(t)$), then $k = 0$ ($\ell = 0$). Contract all edges on the sub-path $P_a \stackrel{\text{def}}{=} (s, a_1, \dots, a_k)$ such that P_a is reduced to an edge (s, v) . Likewise, contract all edges on the sub-path $P_b \stackrel{\text{def}}{=} (b_1, \dots, b_\ell, t)$ such that P_b is reduced to an edge (v, t) . Denote the resulting graph

by G' . Since P is chordless, then there are no edges between (a_i, b_j) for all $i \in [1, k]$ and all $j \in [1, \ell]$. Therefore, following the contraction, s and t are not adjacent in the new graph G' , and hence separable.

Let $S' \in \mathcal{S}_{s,t}(G')$ be a minimal s, t -separator in G' . By construction, $v \in N_{G'}(s) \cap N_{G'}(t)$, and hence $v \in S'$. It is left to show that $S' \in \mathcal{S}_{s,t}(G)$. Let $C_s(G'-S')$ and $C_t(G'-S')$ denote the full connected components of $G'-S'$ containing s and t respectively. Define $D_s(G-S') \stackrel{\text{def}}{=} C_s(G'-S') \cup \{a_1, \dots, a_k\}$ and $D_t(G-S') \stackrel{\text{def}}{=} C_t(G'-S') \cup \{b_1, \dots, b_\ell\}$. By construction, $D_s(G-S')$ and $D_t(G-S')$ are disjoint, non-adjacent, and $G[D_s(G-S')]$ ($G[D_t(G-S')]$) are both connected components in G . Since $C_s(G'-S')$ and $C_t(G'-S')$ are full components of S' in G' , and $D_s(G-S') \supseteq C_s(G'-S')$ and $D_t(G-S') \supseteq C_t(G'-S')$, then $D_s(G-S')$ and $D_t(G-S')$ are full components of S' in G . By Lemma 1, $S' \in \mathcal{S}_{s,t}(G)$. \square

Theorem 18 provides a characterization of when a vertex v is included in a minimal s, t -separator. By reduction from the 3-IN-A-PATH problem we conclude that deciding whether there is a minimal s, t -separator containing a subset $I \subseteq V(G)$ is an NP-complete problem.