All-Subsets Important Separators with Applications to Sample Sets, Balanced Separators and Vertex Sparsifiers in Directed Graphs

Aditya Anand * Euiwoong Lee † Jason Li ‡ Thatchaphol Saranurak§

Abstract

Given a directed graph G with n vertices and m edges, a parameter k and two disjoint subsets $S, T \subseteq V(G)$, we show that the number of *all-subsets important separators*, which is the number of A-B important vertex separators of size at most k over all $A \subseteq S$ and $B \subseteq T$, is at most $\beta(|S|, |T|, k) = 4^k \binom{|S|}{\leq k} \binom{|T|}{\leq 2k}$, where $\binom{x}{\leq c} = \sum_{i=1}^{c} \binom{x}{i}$, and that they can be enumerated in time $\mathcal{O}(\beta(|S|, |T|, k)k^2(m+n))$. This is a generalization of the folklore result stating that the number of A-B important separators for two fixed sets A and B is at most 4^k (first implicitly shown by Chen, Liu and Lu Algorithmica '09). From this result, we obtain the following applications:

1. We give a construction for detection sets and sample sets in directed graphs, generalizing the results of Kleinberg (Internet Mathematics' 03) and Feige and Mahdian (STOC' 06) to directed graphs.
2. Via our new sample sets, we give the first FPT algorithm for finding balanced separators in directed graphs parameterized by k, the size of the separator. Our algorithm runs in time $2^{\mathcal{O}(k)} \cdot (m+n)$.
3. Additionally, we show a $\mathcal{O}(\sqrt{\log k})$ approximation algorithm for finding balanced separators in directed graphs in polynomial time. This improves the best known approximation guarantee of $\mathcal{O}(\sqrt{\log n})$ and matches the known guarantee in undirected graphs by Feige, Hajiaghayi and Lee (SICOMP' 08).
4. Finally, using our algorithm for listing all-subsets important separators, we give a deterministic construction of vertex cut sparsifiers in directed graphs when we are interested in preserving min-cuts of size upto c between bipartitions of the terminal set. Our algorithm constructs a sparsifier of size $\mathcal{O}\left(\binom{t}{\leq 3c} 2^{\mathcal{O}(c)}\right)$ and runs in time $\mathcal{O}\left(\binom{t}{\leq 3c} 2^{\mathcal{O}(c)}(m+n)\right)$, where t is the number of terminals, and the sparsifier additionally preserves the set of important separators of size at most c between bipartitions of the terminals.

*University of Michigan, Ann Arbor

†University of Michigan, Ann Arbor. Supported in part by NSF grant CCF-2236669 and Google

‡Carnegie Mellon University

§University of Michigan, thsa@umich.edu. Supported by NSF Grant CCF-2238138. Partially funded by the Ministry of Education and Science of Bulgaria's support for INSAIT, Sofia University "St. Kliment Ohridski" as part of the Bulgarian National Roadmap for Research Infrastructure.

Contents

1 Introduction

The study of parameterized algorithms, or fixed-parameter tractability (FPT) (see [CFK+15, DF+13]) along with the classical area of approximation algorithms, has emerged as one of the most promising ways to cope with NP-completeness. Given a decision problem Q with input size n, together with a parameter ℓ, a parameterized (or FPT) algorithm is an algorithm running in time $f(\ell)n^{\mathcal{O}(1)}$ that decides Q. Over the last few years, parameterized algorithms have been studied for most well-studied NP-complete problems. A particular focus has been on graph cut optimization problems, including Multiway Cut, Multicut, Minimum Bisection, Feedback Vertex Set, Balanced Separator [Mar06, FM06, CLL+08, MR14, CLP+19, CKL+20, LN22]. One of the most important tools for graph cut problems has been the technique of *important separators* (see [Mar11] for a brief survey), which has formed a building block in parameterized algorithms for many of these and other problems [Mar06, CLL+08, CLL09, CHM13, LM13, LMR+21].

In this work, our contribution is two-fold. First, we show a new structural result on important separators. Next, we show that using this result, combined with other techniques, leads to interesting consequences. We show that one can compute *sample sets* in directed graphs. This in turn allows us to obtain both an FPT algorithm and an improved approximation algorithm for finding small balanced separators in directed graphs. We also show a deterministic construction of vertex cut sparsifiers in directed graphs. All our algorithms are simple and concise modulo standard results in approximation and parameterized algorithms.

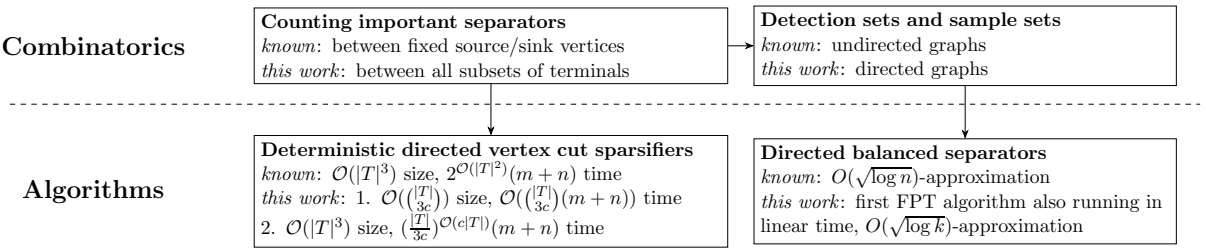

Figure 1: Summary of contribution of our paper

1.1 All-Subsets Important Separators

Important separators have proved to be a very important tool in the design of many FPT algorithms, including fundamental problems such as Multiway Cut, Multicut, Directed Feedback Arc Set. We refer to Chapter 8 of [CFK+15] for various applications of important separators to design parameterized algorithms. Given two disjoint subsets of vertices A, B and another set of vertices X (which may intersect A, B) in a directed graph, we say that X is an A-B separator if there is no directed path from any vertex of A to any vertex of B in $G \setminus X$[1]. For such a separator X, define $R(X)$ to be the set of vertices reachable from A in $G \setminus X$. Then an A-B separator X is called an important A-B separator if it is (a) inclusion-wise minimal - so that for any $v \in X$, $X \setminus \{v\}$ is not an A-B separator and (b) for any other A-B separator $Y \subseteq V(G)$ with $|Y| \leq |X|$, we do not have $R(X) \subset R(Y)$.[2] Informally, important separators are those for which there is no other separator which is further away from A without increasing the cut-size. Marx [Mar06] first introduced the notion of important separators and showed a bound of 4^{k^2} on the number of important A-B separators of size at most k. Later this bound was improved to

[1] $G \setminus X$ denotes the graph obtained from G by deleting the vertex set X along with its incident edges

[2] Throughout the paper, we use the notation $P \subset Q$ to mean that P is a proper subset of Q.

4^k (implicit in the FPT algorithm for Vertex Multiway Cut by Chen, Liu and Lu [CLL09])[3]. Over the last few years, this result has been used extensively to obtain FPT algorithms for various central parameterized problems [MR14, CHM13, CLP+19, LMR+21]. In this work, we show a new structural result on important separators in directed graphs: Given a directed graph G and two disjoint subsets $S, T \subseteq V(G)$, we show that the number of A-B important separators across all $A \subseteq S$ and $B \subseteq T$ is at most $\beta(|S|, |T|, k) = 4^k \binom{|S|}{\leq k}\binom{|T|}{\leq 2k}$, where $\binom{n}{\leq a} = \sum_{i=1}^{a} \binom{n}{i}$. Note that the trivial bound is $4^k 2^{|S|+|T|}$, which follows directly from the fact that there are at most 4^k important A-B separators for any fixed $A \subseteq S, B \subseteq T$. Previously, no such result is known even for undirected graphs.

Throughout this paper, given a graph, we will denote by n the number of vertices and by m the number of edges/arcs.

Theorem 1.1 (All Subsets Important Separators)**.** *Let G be a digraph, k be a positive integer and let $S \subseteq V(G)$ and $T \subseteq V(G)$ be disjoint sets of source and sink vertices. Then there are at most $\beta(|S|, |T|, k) = 4^k \binom{|S|}{\leq k}\binom{|T|}{\leq 2k}$ A-B important separators of size $\leq k$ across all $A \subseteq S$ and $B \subseteq T$. Further, they can be enumerated in time $\mathcal{O}(\beta(|S|, |T|, k) \cdot k^2 \cdot (m + n))$.*

It is not difficult to show that this result is essentially tight: we show this formally in the next lemma, whose proof is deferred to the appendix.

Lemma 1.2. *For any positive integer k, the following two statements hold.*

1. *There exist infinitely many positive integers c, such that for each c, there is a directed graph $G_{c,k}$ and disjoint subsets $S, T \subseteq V(G)$ of vertices with $|T| = c$ and $|S| = 1$ so that there are at least $\binom{|T|}{\leq k}$ A-B important separators of size at most k across all choices $A \subseteq S$, $B \subseteq T$.*
2. *There exist infinitely many positive integers c, such that for each c, there is a directed graph $G_{c,k}$ and disjoint subsets $S, T \subseteq V(G)$ of vertices with $|S| = c$ and $|T| = k + 1$ so that there are at least $\binom{|S|}{\leq k}$ A-B important separators of size at most k across all choices $A \subseteq S$, $B \subseteq T$.*

In the next few subsections, we show that this simple structural result has various interesting consequences, allowing us to obtain many new results which were previously only known for undirected graphs.

1.2 Detection Sets and Sample Sets in Directed Graphs

Kleinberg [Kle03] introduced the concept of a detection set in a graph. The principal motivation for this concept was that of a network failure: given a network, can we compute a small set of representative nodes so that for any small set of failure nodes that cuts communication between two large subsets of nodes, there are two representatives that cannot communicate? Formally, Kleinberg defined an (ϵ, k) detection set for an undirected graph G as a set of terminals $T \subseteq V(G)$ which satisfies the following property. First, define a network failure as a set of vertices X with $|X| \leq k$ so that $G \setminus X$ can be partitioned into $A \cup B$ where $|A|, |B| \geq \epsilon n$ and there are no edges between A and B. Then for every such (vertex) failure set X, the set T must intersect with at least two components of $G \setminus X$. Kleinberg [Kle03] showed that there is a detection set of size $\mathcal{O}(\frac{k^3}{\epsilon} \log \frac{1}{\epsilon})$. For the edge failure version, he showed a bound of $\mathcal{O}(\frac{k}{\epsilon} \log \frac{1}{\epsilon})$ which was subsequently improved to $\mathcal{O}(\frac{k}{\lambda}\frac{1}{\epsilon} \log \frac{1}{\epsilon})$ by Kleinberg et al. [KSS08], where λ is the size of the global minimum cut in the graph. Feige and Mahdian [FM06] showed an improved bound on (vertex) detection sets: they

[3]While these results focused mostly on the undirected version, the same proof works in directed graphs.

show a bound of $\mathcal{O}(\frac{k}{\epsilon}\log\frac{1}{\epsilon})$. Further, they showed a bound of $\mathcal{O}(\frac{k}{\epsilon})$ for the edge failure version, removing the dependence on $\log\frac{1}{\epsilon}$.

Feige and Mahdian [FM06] also studied a strengthening of this notion of detection sets called *sample sets* for undirected graphs: on a high level, these are a small set of terminals T which represent all the small cuts of the graph in proportion to their size, up to some additive error. Concretely, given an undirected graph G and a parameter k, they show that there exists a set of terminals T with $|T| = \mathcal{O}(\frac{k}{\epsilon^2}\log\frac{1}{\epsilon})$, such that for any vertex set X of size $|X| \leq k$ and every connected component C of $G \setminus X$, we have $||C| - \frac{n}{|T|}|C \cap T|| \leq \epsilon n$. Further, the set T can be obtained by simple random sampling, and [FM06] shows that any random subset $T \subseteq V(G)$ of size $\mathcal{O}(\frac{k}{\epsilon^2}\log\frac{1}{\epsilon})$ is a sample set with constant probability.

The most important feature of these results is that the size of the detection set (or sample set) T does not depend on n. It is then natural to ask if such results are possible for directed graphs. Given a digraph G, let us define a network failure as a set of vertices X with $|X| \leq k$, so that $G \setminus X$ can be partitioned into two parts A and B, each of size at least ϵn, such that there is no arc from B to A. The analogous question for detection sets is then: is there a set of terminals T with $|T| = f(k, \epsilon)$, such that for any network failure X there exists a pair $t_1, t_2 \in T$ so that there is no path from t_1 to t_2 in $G \setminus X$? Similarly, the analogous question for sample set becomes: is there a set of terminals T with $|T| = g(k, \epsilon)$ for some function f, so that for any set of vertices X with $|X| \leq k$, and any *strongly connected component (SCC)* C of $G \setminus X$, we have $||C| - \frac{n}{|T|}|C \cap T|| \leq \epsilon n$? We answer both these questions in the affirmative, showing that one can in fact asymptotically match the same bound as in the undirected case, with $f(k, \epsilon) = \mathcal{O}(\frac{k}{\epsilon}\log\frac{1}{\epsilon})$ and $g(k, \epsilon) = \mathcal{O}(\frac{k}{\epsilon^2}\log\frac{1}{\epsilon})$.

Theorem 1.3. *For any directed graph G, given parameters $\epsilon \in (0, 1)$ and $k \in \mathbb{N}$, there is an absolute constant c such that there is an (ϵ, k) detection set of size $f(k, \epsilon) = c\frac{k}{\epsilon}\log\frac{1}{\epsilon}$. Further, a random set of $f(k, \epsilon)$ vertices must be an (ϵ, k) detection set with probability at least $\frac{2}{3}$.*

Theorem 1.4. *For any directed graph G, given parameters $\epsilon \in (0, 1)$ and $k \in \mathbb{N}$, there is an absolute constant c such that there is an (ϵ, k) sample set of size $f(k, \epsilon) = c\frac{k}{\epsilon^2}\log\frac{1}{\epsilon}$. Further, a random set of $f(k, \epsilon)$ vertices must be an (ϵ, k) sample with probability at least $\frac{2}{3}$.*

While we believe that this is of independent interest, we also show a similar connection to parameterized and approximation algorithms for finding balanced cuts in directed graphs along the lines of the undirected case as in [FM06], as discussed in the following subsections. In fact, by using a slightly more nuanced analysis, our results generalize that of [FM06] even for undirected graphs.

Finally, we note that one can easily prove that there exists an absolute constant c such that a random subset of size $c\frac{k\log n}{\epsilon^2}$ is a sample set with constant probability in a directed graph. This follows from a simple application of Chernoff bounds and a union bound noting that the number of vertex sets of size at most k is at most $\binom{n}{k}$ ($\binom{m}{k}$ in the case of edge sets). However, for most applications this bound is too weak. For instance, our parameterized algorithm for Directed Balanced Separator has an exponential dependency in the size of the sample set, and hence to show results parameterized by k, it is essential that the size of the sample set does not depend on n.

1.3 Directed Balanced Cuts

One of the most well-studied graph partitioning problems, both from the parameterized and approximation algorithms point of view, is the problem of Minimum Bisection. Given an undirected graph G and a parameter k, the goal of the Minimum Bisection problem is to obtain a partition

of the graph into two equal sized parts, such that the number of cut edges is at most k. Minimum Bisection was shown to be fixed-parameter tractable by Cygan et al. [CLP+19] and the current best parameterized algorithm is due to [CKL+20] who show an algorithm with running time $2^{\mathcal{O}(k \log k)} n^{\mathcal{O}(1)}$. Räcke [Räc08] gave an $\mathcal{O}(\log n)$ approximation for (the optimization version of) Minimum Bisection. If only an approximate bisection where both sides have $\Omega(n)$ vertices is desired, then this problem is essentially the Balanced Separator problem which is known to have an FPT algorithm running in time $2^{\mathcal{O}(k)}(m+n)$ due to Feige and Mahdian [FM06] and an $\mathcal{O}(\sqrt{\log n})$ approximation algorithm in polynomial time using the seminal result of Arora, Rao and Vazirani [ARV09]. Feige et al. [FHL08] showed that in fact this guarantee can be made $\mathcal{O}(\sqrt{\log k})$, and also showed that one can compute vertex separators with the same approximation ratio. In directed graphs, the Minimum Bisection problem has been typically studied as Directed Bisection: Given a directed graph G with even number of vertices, is it possible to partition the vertex set into two equal parts A and B so that the number of arcs from A to B is at most k? This question was first raised by Feige and Yahalom [FY03]. When $k = 0$, they referred to this problem as Oneway Cuts, and showed that even this problem is NP-hard. However, given a Oneway Cuts instance which admits a solution, if one relaxes the requirement so that the algorithm can output a partition (A', B') of $V(G)$ such that there are no arcs directed from A' to B' and $||A'| - |B'|| \leq \epsilon n$ where $\epsilon = \Omega(\frac{1}{\log n})$, the problem now becomes tractable, and they show a polynomial time algorithm for Oneway Cuts. Madathil et al. [MSZ21] showed that the Directed Bisection problem is FPT with respect to k, even when one requires $|A| = |B|$ exactly, on a subclass of directed graphs called semi-complete digraphs, which are the class of directed graphs where for every pair of vertices u and v, there is an arc from u to v or an arc from v to u. In terms of approximation algorithms, Agarwal et al. [ACMM05] showed an $\mathcal{O}(\sqrt{\log n})$ approximation for Directed Balanced Separator , the directed analogue of Balanced Separator, while Even et al. [ENRS99] showed an $\mathcal{O}(\log k)$ approximation, which is the best known approximation guarantee depending only on k.

However, there is no prior work on the fixed-parameter tractability of Directed Balanced Separator or Directed Bisection for general k on general directed graphs, even when we relax the requirement of finding a bisection to that of finding an *approximate bisection*, that is, find a partition (A', B') with $|A'|, |B'| = \Omega(n)$ so that the number of arcs from A' to B' is at most k.

Our result makes the first progress on this problem. Before we state our results, we define the Directed Balanced Separator problem formally. Since all our results work for both the vertex and edge versions, we state them together as one problem. We adapt our definition from the definition of Balanced Separator in [FM06], whose results we generalize. For completeness, we recall the definition of Balanced Separator in [FM06].

Balanced Separator	**Parameter:** k, b
Input: Undirected graph $G = (V, E)$	
Question: Is there a set of vertices (edges) F with $\lvert F\rvert \leq k$, so that in $G \setminus F$, every connected component has size at most bn?	

Directed Balanced Separator	**Parameter:** k, b
Input: Directed graph $G = (V, E)$	
Question: Is there a set of vertices (arcs) F with $\lvert F\rvert \leq k$, so that in $G \setminus F$, every strongly connected component has size at most bn?	

For the sake of clarity and comparison, we state the main result of [FM06]. Given an undirected

graph G, we say that a set of vertices/edges F is a *b-balanced separator* if every connected component of $G \setminus F$ has size at most bn.

Theorem 1.5 ([FM06]). *Given an instance of* Balanced Separator *with $\frac{2}{3} \leq b \leq 1$ there is a randomized algorithm, that for any $\epsilon > 0$, runs in time $2^{\mathcal{O}\left(k \log\left(\frac{1}{\epsilon}\right)/\epsilon^2\right)}(m+n)$ and with constant probability outputs either (a) a set of vertices (edges) F' of size at most k such that every connected component of $G \setminus F'$ has size at most $(b+\epsilon)n$ or (b) concludes correctly that there is no b-balanced separator of size at most k.*

The following theorem, which directly generalizes the result of Theorem 1.5 is our main result. Given a directed graph G, we say that a set of vertices/arcs F is a *b-balanced separator* if every *strongly connected component* (SCC) of $G \setminus F$ has size at most bn.

Theorem 1.6. *Given an instance of* Directed Balanced Separator *there is a randomized algorithm, that for any $\epsilon > 0$, runs in time $2^{\mathcal{O}\left(k \min\{\log \frac{1}{b}, \log k\} \log \frac{1}{\epsilon}/\epsilon^2\right)}(m+n)$ and with constant probability outputs either (a) a set of vertices (arcs) F' of size at most k such that every strongly connected component of $G \setminus F'$ has size at most $(b+\epsilon)n$ or (b) concludes correctly that there is no b-balanced separator of size at most k.*

We observe that our algorithm has a running time of $2^{\mathcal{O}\left(k \log\left(\frac{1}{\epsilon}\right)/\epsilon^2\right)}(m+n)$ for any $b = \Omega(1)$, matching the run-time of Theorem 1.5 for $\frac{2}{3} \leq b \leq 1$ while also extending to any parameter $b \in (0,1)$. We also observe that Theorem 1.6 implies Theorem 1.5, since given any undirected graph G, one can create a directed graph H on the same vertex set, so that for every edge $\{u,v\} \in E(G)$, we have the two arcs $(u,v),(v,u) \in E(H)$ and we can apply Theorem 1.6 to obtain Theorem 1.5.

Our algorithm for Directed Balanced Separator can be used to solve (approximate) Directed Bisection as well. To see this, given a graph G, observe that in any bisection (A,B) so that the number of arcs going from A to B is at most k, the set F of these at most k arcs form a $\frac{1}{2}$-balanced separator, so that in $G \setminus F$, every strongly connected component is of size at most $\frac{n}{2}$. Using Theorem 1.6 with $b = \frac{1}{2}$, we can find a set of arcs F' with $|F'| \leq k$ that forms a $(\frac{1}{2}+\epsilon)$ balanced separator. Finally, note that the strongly connected components of the graph $G \setminus F'$ form a Directed Acyclic Graph (DAG), and each strongly connected component of $G \setminus F'$ has size at most $(\frac{1}{2}+\epsilon)n$. It follows that there is a topological ordering $\{C_1, C_2 \ldots C_\ell\}$ of these strongly connected components of $G \setminus F'$, such that there is no arc from C_j to C_i for $i,j \in [\ell]$ with $j > i$. Therefore we can pick some prefix of strongly connected components in the topological ordering of $G \setminus F'$ to obtain a set A', so that both $|A'|, |V(G) \setminus A'| \geq (1-2\epsilon)\frac{n}{4}$, and in $G \setminus F'$, there are no arcs from $V(G) \setminus A'$ to A'. It follows that the set of arcs F' forms an (approximate) directed bisection. Note that the approximation is only in the balance, not in the number of arcs cut.

Next, we show a $\mathcal{O}(\sqrt{\log k})$ approximation for Directed Balanced Separator. This improves both the $\mathcal{O}(\sqrt{\log n})$ approximation of [ACMM05] and the $\mathcal{O}(\log k)$ approximation given by Even et al. [ENRS99] for approximating Directed Balanced Separator in polynomial time.

Theorem 1.7. *There is an $\mathcal{O}(\sqrt{\log k})$ approximation to* Directed Balanced Separator *in polynomial time. Formally, given an instance of* Directed Balanced Separator *with $b = \Omega(1)$, suppose there is a set of vertices (arcs) F with $|F| \leq k$, so that every strongly connected component of $G \setminus F$ has at most bn vertices. Then there is a polynomial time randomized algorithm that with constant probability finds a set of vertices (arcs) F' with $|F'| \leq \mathcal{O}(k\sqrt{\log k})$ so that in $G \setminus F'$, every strongly connected component has size at most $b'n$ for some $b' < 1$ depending on b.*

Note that for this theorem, we do not optimize b', unlike our FPT result where we were able to show that $b' \leq b + \epsilon$ for some suitably chosen ϵ. Still, we remark that is not a limitation of our framework, but inherent in [ACMM05] due to the use of the ARV separation theorem [ARV09].

1.4 Vertex Cut Sparsifiers

Vertex sparsification is a fundamental problem in various settings. Broadly, given a graph and a terminal set T, a vertex sparsifier is a smaller graph (with size typically depending only on $|T|$) that preserves some (cut based) property of the terminals T. Moitra [Moi09] first introduced a version of vertex sparsification in undirected graphs. Chuzhoy [Chu12] generalized this notion by allowing Steiner nodes in the sparsifier. Both these notions preserve edge cuts between bi-partitions of the terminal set.

For our setting, we focus on directed graphs and vertex cuts. Given a directed graph G and a set of terminals T along with an integer c, a (c, T) vertex cut sparsifier for G is another graph G' with $T \subseteq V(G')$, so that for every partition $\{A, B\}$ of T (so that $A \cup B = T$ and $A \cap B = \emptyset$), if the size of an A-B vertex min-cut is at most c in G, then the size of an A-B vertex min-cut is the same in both in G and G'. In other words, the goal is to find a vertex sparsifier that preserves vertex min-cuts of size at most c between sets of terminals. Here we allow the deletion of terminals as well.

Kratsch and Wahlström [KW12, KW20] first studied this notion of vertex sparsification for vertex cuts (without the parameter c) and showed that given G, T, there is a randomized polynomial time algorithm that obtains a $(|T|, T)$ vertex sparsifier G' for G with $|V(G')| \leq \mathcal{O}(|T|^3)$. This bound was improved to $\mathcal{O}(|T|^2)$ by [HLW21] for the special case of directed acyclic graphs. Their algorithm runs in linear time for fixed $|T|$, but is still randomized as it needs to compute a representation of gammoids.

One can then ask the question as to what is known about deterministic algorithms. Recently, Misra et al. [MPRS20] showed that a representation of a gammoid with rank r over a ground set of m elements can be constructed in time $\mathcal{O}(\binom{m}{r} m^{\mathcal{O}(1)})$ deterministic time. The technique of [KW12] needs to compute a representation of a gammoid whose ground set has size n, the number of vertices and rank $|T|$, the number of terminals. We therefore obtain the following result.

Theorem 1.8 ([KW12], [MPRS20])**.** *Given a n-vertex graph G and a terminal set T, there is a deterministic algorithm that runs in time $\mathcal{O}(\binom{n}{|T|} n^{\mathcal{O}(1)})$ and computes a $(|T|, T)$ vertex cut sparsifier for G of size at most $\mathcal{O}(|T|^3)$.*

However, this algorithm (which is an XP algorithm in the notation of parameterized complexity) can be easily improved to a (still deterministic) FPT algorithm. The reason for this is simple: there are at most $2^{|T|}$ partitions of the terminal set T. For each partition $A \cup B$, we can compute an (A, B) directed min-vertex cut M. Note that this min-cut is of size at most $|T|$, since one can simply delete all the terminals to obtain a cut. This gives a set X of $2^{|T|} \cdot |T|$ vertices. It can now be shown that we can apply the closure operation to the set $V(G) \setminus X$, where we simply delete all vertices of $V(G) \setminus X$, and for each vertex pair (u, w) such that there are vertices (v_1, v_2) with $v_1, v_2 \in V(G) \setminus X$ with $(u, v_1), (v_2, w) \in E(G)$ we add an arc (u, w). This results in a sparsifier of size $\mathcal{O}(|T|2^{|T|})$. One can now set $n = \mathcal{O}(|T|2^{|T|})$ in Theorem 1.8 to obtain a sparsifer of size $2^{\mathcal{O}(|T|^2)}$ in time $2^{\mathcal{O}(|T|^2)}$. The total running time is $2^{\mathcal{O}(T)}(m + n) + 2^{\mathcal{O}(|T|^2)}$.

A similar line of research considers edge and vertex cuts in undirected graphs. For edge cuts in undirected graphs, Chalermsook et al. [CDK+21] showed an upper bound of $\mathcal{O}(|T|c^4)$ which was later improved to $\mathcal{O}(|T|c^3)$ by Liu [Liu23]. Both these algorithms are randomized. Saranurak and Yingchareonthawornchai [SY22] considered the vertex version in undirected graphs, and gave a deterministic algorithm to compute a sparsifier of size $\mathcal{O}(|T|2^{\mathcal{O}(c^2)})$ in time $\mathcal{O}(m^{1+o(1)}2^{\mathcal{O}(c^2)})$.

However, no improvement on the FPT algorithm with running time $2^{O(|T|)}(m + n) + 2^{O(|T|^2)}$ based on Theorem 1.8 is known for deterministic algorithms for vertex cut sparsifiers in directed

graphs. We show the following result, which shows that if one is interested in preserving cuts of size at most c, then a better deterministic algorithm is possible.

Theorem 1.9. *Given a directed graph G, a set of terminals T and integer c, there is a deterministic algorithm that runs in time $\mathcal{O}(\psi(|T|,c)\cdot(m+n))$ and computes a (c,T) vertex sparsifier G' for G of size at most $\psi(|T|,c)$, where $\psi(|T|,c)=\binom{|T|}{\leq 3c}2^{\mathcal{O}(c)}$. Additionally, G' satisfies $T\subseteq V(G')\subseteq V(G)$, with the property that for every partition $A\cup B$ of T and every subset $X\subseteq V(G)$ with $|X|\leq c$, X is an A-B important separator in G if and only if $X\subseteq V(G')$ and X is an A-B important separator in G'.*

While the dependence on c in the exponent for the size of the sparsifier seems undesirable, Theorem 1.9 shows that the vertex sparsifier G' constructed by our algorithm is more powerful: it in-fact "preserves" all A-B important separators for any partition (A,B) of T. Additionally, if one wants only a vertex sparsifier, we note that simply running the algorithm in Theorem 1.8 after applying our result in Theorem 1.9 gives a vertex sparsifier of size $\mathcal{O}(|T|^3)$, but now the running time is $\left(\frac{|T|}{c}\right)^{\mathcal{O}(c|T|)}$. Thus, together, we obtain a vertex cut sparsifier of size $\mathcal{O}(|T|^3)$ in time $\binom{|T|}{\leq 3c}2^{\mathcal{O}(c)}(m+n)+\left(\frac{|T|}{c}\right)^{\mathcal{O}(c|T|)}$.

Corollary 1.10. *Given a graph G and a terminal set $T\subseteq V(G)$, there is a deterministic algorithm that runs in time $\binom{|T|}{\leq 3c}2^{\mathcal{O}(c)}(m+n)+\left(\frac{|T|}{c}\right)^{\mathcal{O}(c|T|)}$ and obtains a (c,T) vertex cut sparisifer for G with $\mathcal{O}(|T|^3)$ vertices.*

1.5 Techniques and Overview

In this section we give a high-level overview of our techniques. For the sake of simplicity and clarity we sometimes omit small details in this section.

All-Subset Important separators: Our result on counting all-subset important separators is obtained using the connection between important separators and *closest sets*. For simplicity, in this overview, suppose that we are given a single source s and a set of target vertices T. Our goal in this simplified setting shall be to bound the number of (s,B) important separators of size exactly k across all $B\subseteq T$, and we will show a weaker bound of $4^k\binom{|T|}{\leq (k+1)^2}$. Before we describe the main idea, we remark that our eventual goal will be to show that every (s,B) important separator of size k for some $B\subseteq T$ is an (s,B') important separator for some $B'\subseteq B$ with $|B'|\leq (k+1)^2$. It is easy to see that this suffices, since there are only $\binom{|T|}{\leq (k+1)^2}$ choices for B', and there are at most 4^k such (s,B') important separators of size k for a fixed choice of B'.

Fix $B\subseteq T$, and consider an (s,B) important separator X of size k. Consider the (vertex) min-cut X' between X and B[4]. If this min-cut X' has size strictly less than $|X|=k$, X cannot be important: X' would allow more vertices to be reachable from s while having a size strictly less than k. This means that X by itself must be a (X,B) min-cut. Using simple cut-flow duality, it must be the case that there are $|X|=k$ vertex disjoint paths from X to (some) vertices in B.

However, this flow property can be strengthened: it is easy to see that we can in fact assume that every such important separator X is the (X,B) min-cut "closest" to B. (We say X is closest to B if X is the unique (X,B) min-cut.) But it can be shown that closest sets have a stronger flow

[4]Throughout this paper, we allow deletion of source and sink vertices in vertex cuts

property: we can now conclude that for every vertex $v \in X$, there are $k+1$ paths from X to B, which are vertex disjoint except that two of these paths both start at v (Lemma 2.2).

Why does this help? Fix a $v \in X$, and fix the $k+1$ paths from X to B. Suppose the $k+1$ endpoints of these paths are $B_v^* \subseteq B$. Consider any set of vertices X' that is a (s, B_v^*) separator of size $\leq k$ that is "closer" to B_v^* than X. (Here "closer" means that the set of vertices reachable from s in $G \setminus X'$ is a superset of that in $G \setminus X$). Then X' must delete v! For if not, since there are $k+1$ paths from X to B_v^* which are vertex disjoint except 2 paths which begin at v, and $|X'| \leq k$, X' cannot be a (s, B_v^*) separator.

Applying the above argument for every $v \in X$ and letting $B' := \cup_{v \in X} B_v^*$ allows us to conclude that there is no (s, B') separator X' that is closer to B' than X. Thus, in fact, among all the sets of size at most k separating s from B', X is a set which is closest to B'. This means X must be an important (s, B') separator. But $|B'| = |\cup_{v \in X} B_v^*| \leq k(k+1) \leq (k+1)^2$, and hence we are done. Our actual result uses a slightly more subtle argument using augmenting paths to obtain such a B' with $|B'| \leq 2k$.

Here we remark that Lemma 4.10 of Lokshtanov et al. [LRS+25] showed a similar result: it claims that there is a set $B' \subseteq B$ with $|B'| \leq k+1$ such that any (s, B) important separator is an (s, B') important separator. However, there is a gap in the proof which leads to an unavoidable factor of 2 loss. Indeed, as we show in Lemma A.1, there is a graph G, source vertex $s \in V(G)$ and $B \subseteq V(G)$ for which $|B'| = 2k$ is necessary, and hence our result is tight. Our high-level idea is flow-based, similar to [LRS+25]. But as described above, the direct application of this flow-based idea gives a bound of $k(k+1)$, which we refine using an augmenting path based argument to obtain the tight bound of $2k$.

Detection sets and Sample sets for Directed Graphs: To obtain detection sets and sample sets for directed graphs, on a high level, we follow the approach of Feige and Mahdian [FM06]. We focus on sample sets in this brief overview. First, we consider the family of sets that arise as a strongly connected component after deleting an arbitrary k vertices in the given digraph. Our key contribution is to show that this family has VC-dimension $\mathcal{O}(k)$. Using the standard results on ϵ-samples, similar to that in [FM06], it then follows that a random set of $\mathcal{O}(\frac{k}{\epsilon^2} \log \frac{1}{\epsilon})$ vertices is a sample set with constant probability. Thus the question becomes: how can we bound the VC-dimension of the family of sets $\mathcal{S}$ formed by strongly connected components after deleting some k vertices? We note that the techniques in [FM06] to bound the VC-dimension do not extend to directed graphs, and hence a different approach is needed. We accomplish this by showing a connection to our result on all-subset important separators.

Recall the definition of VC-dimension: given a set family $\mathcal{S}$ on a universe U, the VC-dimension is the size of the largest set $U' \subseteq U$ shattered by $\mathcal{S}$. Here U' is shattered by $\mathcal{S}$ if for every subset $U'' \subseteq U'$ there exists an $S \in \mathcal{S}$ with $|U' \cap S| = U''$. Thus to prove a bound of $\mathcal{O}(k)$ for VC-dimension for our case, it is enough to show that there exists some constant c so that any set of terminals $T \subseteq V(G)$ with $|T| \geq ck$ cannot be shattered by the set family $\mathcal{S}$ which consists of the strongly connected components C in $G \setminus F$ across all sets $F \subseteq V(G)$ with $|F| \leq k$. Fix any T of size $\geq ck$ for large enough c which we choose later. For the sake of contradiction, assume that the set T can be shattered. Fix a "pattern" $P \subseteq T$, $P \neq \emptyset$. Since T can be shattered by $\mathcal{S}$, there exists a set F of at most k vertices, so that in $G \setminus F$, there is a strongly connected component C that satisfies $C \cap T = P$. Fix some terminal $t \in P$. Then it must be the case that P is the exact set of terminals that can both reach t and can be reached from t in $G \setminus F$. Equivalently, it is the exact set of terminals that can be reached from t in both $G \setminus F$ and $G^R \setminus F$, where G^R is obtained by reversing every arc of G. This motivates us to define Reach(t) for each $t \in T$, the collection of subsets of

terminals that can be reached from t after deleting a set F of size at most k. Formally,

$$\text{Reach}(t) = \{Q \subseteq T \mid \text{there exists } F \subseteq V(G), |F| \leq k \text{ such that } Q \text{ is reachable from } t \text{ while } T \setminus Q \text{ is unreachable from } t \text{ in } G \setminus F\}$$

If we can bound $|\text{Reach}(t)|$ for any graph by a function $B(|T|, k)$, then applying this result twice, once on G and once on G^R, will bound the number of such patterns P with $t \in P$ by $(B(|T|, k))^2$. Then a simple union bound over all $t \in T$ will give that the number of patterns $P \subseteq T$ is at most $|T|B(|T|, k))^2$. If this is less than $2^{|T|}$ whenever $|T| \geq ck$, then we have a contradiction to the fact that T can be shattered. Thus the key question is to obtain $B(|T|, k)$, a good bound on $|\text{Reach}(t)|$.

We do this as follows. Suppose $T' \in \text{Reach}(t)$ some $T' \subseteq T$. Then there exists a set F of at most k vertices after whose deletion the set of reachable terminals from t is exactly T'. We then show that there is a $(t, T \setminus T')$ *important separator* X of size at most k whose deletion gives the same reachability set T' as that after deleting F. But now using our bound on important separators, we can show that $B(|T|, k) \leq 4^k \binom{|T|}{\leq 2k}$. It is then easy to see that $|T|B(|T|, k)^2 < 2^{|T|}$ whenever $|T| \geq ck$ for some large enough (absolute) constant c.

FPT algorithm for Directed Balanced Separators: To obtain the algorithm for Directed Balanced Separator, we follow the approach of [FM06] in the undirected case. For simplicity, in this overview, we work with vertex separators, assume that $b = \frac{1}{2}$, and we will only return a $b' = (\frac{3}{4} + 2\epsilon)$-balanced separator (as opposed to the $(b + \epsilon)$ balance guaranteed in our result). We are given that there is a set of vertices F with $|F| \leq k$, so that every strongly connected component of $G \setminus F$ has at most $\frac{1}{2}n$ vertices. We first compute an (ϵ, k) sample set T of size $\mathcal{O}\left(\frac{k \log \frac{1}{\epsilon}}{\epsilon^2}\right)$ using our result on sample sets, Theorem 1.4. Since T is a sample set, each SCC of $G \setminus F$ has at most $(\frac{1}{2} + \epsilon)|T|$ terminals. We will try to find a set of vertices F' such that $|F'| \leq k$ and every SCC of $G \setminus F'$ has at most $(\frac{3}{4} + \epsilon)|T|$ terminals. The property of sample sets will then again imply that F' is a $(\frac{3}{4} + 2\epsilon)$-balanced separator.

Consider the toplogical ordering $C_1, C_2 \ldots C_\ell$ of the SCC's of $G \setminus F$, so that there is no arc from C_j to C_i for $i, j \in [\ell]$ with $j > i$. Consider the smallest index i^* so that the union $L = \bigcup_{i=1}^{i^*} C_i$ contains at least $(\frac{1}{4} - \epsilon)|T|$ terminals. Since no component contains more than $(\frac{1}{2} + \epsilon)|T|$ terminals, it follows that L contains at most $\frac{3}{4}|T|$ terminals. Also, by definition, $V(G) \setminus L$ contains at most $(\frac{3}{4} + \epsilon)|T|$ terminals.

The algorithm proceeds as follows. First, guess $T \cap L$ and $T \cap (V(G) \setminus L)$. This can be done in time $2^{\mathcal{O}\left(\frac{k \log \frac{1}{\epsilon}}{\epsilon^2}\right)}$ since $|T| = \mathcal{O}\left(\frac{k \log \frac{1}{\epsilon}}{\epsilon^2}\right)$. Next, we compute a directed min-vertex cut between $T \cap (V(G) \setminus L)$ and $T \cap L$. Since F is a vertex cut between these sets of size at most k, this min-cut F' must be of size at most k as well. But then the set of vertices F' is a $(\frac{3}{4} + \epsilon)$-balanced separator with respect to the set T. Using the property of sample sets, it follows that F' is a $(\frac{3}{4} + 2\epsilon)$-balanced separator in G.

In order to make the loss in the balance factor only an additive $\mathcal{O}(\epsilon)$ we need a slightly more sophisticated argument. We accomplish this using a reduction to the Skew Separator problem, where we are given pairs $(s_i, t_i), i \in [\ell]$, and the goal is to separate s_i from all t_j, $j \leq i$, for each i, by deleting minimum number of vertices. This problem, which is a special case of Directed Multicut, was first defined by [CLL+08] in their FPT algorithm for Directed Feedback Arc Set.

Approximation algorithm for Directed Balanced Separators: Our $\mathcal{O}(\sqrt{\log k})$ approximation algorithm proceeds similar to the FPT algorithm. The only key technical difference is that to

compute a balanced separator with respect to the terminal set, we cannot guess where the terminals of the sample set are (this requires FPT time); instead we use the algorithm of [ACMM05] to obtain a balanced separator with respect to the sample set. While the algorithm of [ACMM05] has an approximation ratio of $\mathcal{O}(\sqrt{\log n})$, when we need a balanced separator with respect to a terminal set T, this algorithm can be made to work with an approximation ratio of $\mathcal{O}(\sqrt{\log |T|})$. On a high level, the reason is quite simple: the structure theorem of Arora, Rao and Vazirani [ARV09] is a statement about vectors. As long as we apply the structure theorem to only $|T|$ vectors, we get an approximation ratio of $\mathcal{O}(\sqrt{\log |T|})$.

Vertex Sparsifiers: Our result on vertex sparsifiers is similar in spirit to the result by Kratsch and Wahlström [KW12] and proceeds by identifying *irrelevant vertices*. However instead of using techniques based on matroids, we use our result on all-subsets important separators. Given a terminal set T, we are interested in preserving cuts across partitions of the terminal set of size up to c. We will define an irrelevant vertex as a vertex that is (a) not in T and (b) not in any A-B important separator of size $\leq c$ for any partition $A \cup B$ of T. By our result on all-subsets important separators, we can bound the number of relevant (non-irrelevant) vertices by $\binom{|T|}{\leq 3c}2^{\mathcal{O}(c)}$, and identify this set in time $\binom{|T|}{\leq 3c}2^{\mathcal{O}(c)}(m+n)$. Thus at least $|V(G)| - |T| - \binom{|T|}{\leq 3c}2^{\mathcal{O}(c)}$ vertices are irrelevant. Let I be the set of irrelevant vertices. We now apply the standard operation of "closing" the irrelevant vertices I, which is to delete each $v \in I$, and for every pair (w, y) where w is an in-neighbour of v_1 and y is an out-neighbour of v_2 for some $v_1, v_2 \in I$, to add the arc (w, y). This operation is the vertex equivalent of edge contraction. The key observation here is that the closure operation *can be applied to the entire set I at once* as opposed to applying it to one vertex of I and restarting the whole algorithm. This helps us achieve a linear-time algorithm. Thus we obtain another graph G' with $V(G') \subseteq V(G)$, where $|V(G')| \leq |V(G) \setminus I| \leq \binom{|T|}{\leq 3c}2^{\mathcal{O}(c)}$.

2 All-Subsets Important Separators

The goal of this section is to prove our result on all-subsets important separators.

Theorem 1.1 (All Subsets Important Separators)**.** *Let G be a digraph, k be a positive integer and let $S \subseteq V(G)$ and $T \subseteq V(G)$ be disjoint sets of source and sink vertices. Then there are at most $\beta(|S|, |T|, k) = 4^k\binom{|S|}{\leq k}\binom{|T|}{\leq 2k}$ A-B important separators of size $\leq k$ across all $A \subseteq S$ and $B \subseteq T$. Further, they can be enumerated in time $\mathcal{O}(\beta(|S|, |T|, k) \cdot k^2 \cdot (m+n))$.*

As described in the introduction, our proof proceeds by capturing the relationship between important separators and closest sets. We then use well-known connectivity properties of closest sets to obtain our result.

Definition 2.1 (Closest set)**.** *For sets of vertices $X, T \subseteq V(G)$, X is closest to T if X is the unique vertex mincut between X and T.*

We use the following well-known fact about closest cuts.

Lemma 2.2 (Lemma 18 of [HLW21])**.** *X is closest to T if and only if for each vertex $v \in X$, there are $|X| + 1$ paths from X to T that are vertex-disjoint except at v, which appears in exactly two of these paths.*

Definition 2.3 (Minimal separators)**.** *Given $X, S, T \subseteq V(G)$ such that $S \cap T = \emptyset$, X is called an S-T (inclusion-wise) minimal separator if in $G \setminus X$, there is no path from any $s \in S$ to any $t \in T$, and further, for every $v \in X$, there is a path between some $s \in S$ and $t \in T$ in $G \setminus (X \setminus \{v\})$.*

Definition 2.4. *Given $X, S \subseteq V(G)$, the reachability set of the source vertices S after removing X, that is, the set of all vertices v reachable from some vertex of S via a directed path after removing X, is denoted by $R_S^G(X)$.*

Definition 2.5. *Given $X, S, T \subseteq V(G)$ with $S \cap T = \emptyset$, a minimal S-T separator X is called an S-T important separator if for every S-T separator $X' \subseteq V(G)$ with $|X'| \leq |X|$, $R_S^G(X') \supset R_S^G(X)$ does not hold.*

We will drop the subscript or the superscript in $R_S^G(X)$ when the graph or the set of source vertices is clear from the context.

Informally, the idea of important separators is similar to that of a closest set: they are separators which cannot be pushed "closer" to the sink, without increasing the cut-size. The next lemma captures this relationship formally.

Lemma 2.6. *Given $X, S, T \subseteq V(G)$, X is an important S-T separator if and only if (a) X is a minimal S-T separator and (b) X is closest to T.*

Proof. Notice that by definition, any important S-T separator must be a minimal S-T separator. First, suppose to the contrary that X is an important S-T separator but X is not closest to T. Let Y be a vertex min-cut between X and T with $Y \neq X$. Since X is not closest to T, such a Y indeed does exist. Clearly, $|Y| \leq |X|$. We claim that $R(X) \subset R(Y)$, which will contradict that X is an important separator. First, let us show that $R(X) \subseteq R(Y)$. Suppose for some vertex v, $v \in R(X)$ but $v \notin R(Y)$. Since Y is a min-cut, there are vertex disjoint paths between each vertex $y \in Y$ and some vertex $t(y) \in T$ such that no internal vertex of these paths is from X. Consider a path P from some vertex $s \in S$ to v in $G \setminus X$. There must be such a path, since $v \in R(X)$. Let y be a vertex of Y on this path: there has to be such a vertex y since $v \notin R(Y)$. Then y is reachable from s in $G \setminus X$. Also, there is a path from y to $t(y)$ which does not contain any vertex of X. It follows that there is a path from s to $t(y)$ in $G \setminus X$, which is a contradiction since X is an S-T separator. Since X is a minimal S-T separator and $Y \neq X$, it follows that there exists a $v \in X \cap R(Y)$, and hence $R(X) \subset R(Y)$.

Next, suppose that X is both closest to T and is also a minimal S-T separator. Assume for the sake of contradiction that X is not important. Since X is not important, there is another separator X' with $|X'| \leq |X|$ which is a minimal $S-T$ separator, and further $R(X) \subset R(X')$. Since $X' \neq X$, there exists a $v \in X \cap (V(G) \setminus X')$. By Lemma 2.2, for every vertex $v \in X$ there exists a collection of $|X| + 1$ paths from X to T which are vertex disjoint except at v, which appears twice. Since X' does not delete v, it follows that there is at least one such path P from some vertex $x \in X$ to some vertex $t \in T$ which survives in $G \setminus X'$. We next show that there must be a path from some vertex of S to x in $G \setminus X'$ and this will imply a path from S to t in $G \setminus X'$, a contradiction.

Consider any path from some vertex $s \in S$ to x, all of whose internal vertices are not in X. Such a path must exist, for otherwise X will not be a minimal S-T separator since $X \setminus \{x\}$ will also be an S-T separator. Every internal vertex on this path is reachable from s in $G \setminus X$. But $R(X) \subseteq R(X')$, which means every internal vertex on this path is reachable from s in $G \setminus X'$ as well. This in turn means that there is a path from s to x in $G \setminus X'$, a contradiction. □

Lemma 2.7 (Theorem 8.51 of [CFK+15])**.** *Given a digraph G and two disjoint sets $A, B \subseteq V(G)$, there are at most 4^k A-B important separators of size $\leq k$.*

Equipped with these results, we are now ready to prove our main theorem in this section which bounds the number of important separators between across all subsets $A \subseteq S$ and $B \subseteq T$.

Theorem 1.1 (All Subsets Important Separators)**.** *Let G be a digraph, k be a positive integer and let $S \subseteq V(G)$ and $T \subseteq V(G)$ be disjoint sets of source and sink vertices. Then there are at most $\beta(|S|,|T|,k) = 4^k \binom{|S|}{\leq k}\binom{|T|}{\leq 2k}$ A-B important separators of size $\leq k$ across all $A \subseteq S$ and $B \subseteq T$. Further, they can be enumerated in time $\mathcal{O}(\beta(|S|,|T|,k) \cdot k^2 \cdot (m+n))$.*

Proof. Let X be an important A-B separator for some $A \subseteq S$ and $B \subseteq T$ with $|X| \leq k$. We will show that S is also an A^*-B^* important separator for some $A^* \subseteq S$ and $B^* \subseteq T$ with $|A^*| \leq k$ and $|B^*| \leq 2k$.

By Lemma 2.6, it must be the case that X is closest to B. Also, there are vertex disjoint paths from each vertex x of X to some vertex $t(x)$ of B. Also, we know that for every v there exists $|X|+1$ disjoint paths from X to B which are vertex disjoint except at v which appears twice.

Consider the vertex-flow network H obtained by adding a super-source s^* and a super-sink t^*. We add outgoing arcs from s^* to each vertex of X, and incoming arcs to t^* from each vertex of B. The capacity of each vertex is 1. Let f be the (maximum) flow corresponding to any set of $|X|$ vertex disjoint paths between X and B, and let the endpoints of these paths in B be the set Y. Now we do the following operation for each vertex v. Consider the modified flow network H_v obtained from H where the underlying graph is same, but the vertex capacity of v is raised to 2. f is not a maximum flow in H_v, since there exists $|X|+1$ paths from X to B which are vertex disjoint except at v which appears in two of these paths. Thus, there is an augmenting path for f in H_v. We augment the flow f along this augmenting path to obtain a new flow f_v. The new flow has one additional endpoint in B, so that the endpoints of the paths of this new flow f_v are either in Y or at the newly added vertex $u \in B \setminus Y$. Let us denote by Y_v the set of the new endpoints.

Let $B^* = \bigcup_{v \in X} Y_v$. Notice that $|B^*| \leq 2k$, since each $Y_v = Y \cup \{u\}$ for some vertex u for each $v \in X$. Also since X is minimal, there must exist paths P_v for each $v \in X$ from some vertex $s(v)$ of A to v which do not contain any other vertex of X. Let $A^* = \bigcup_{v \in X} s(v)$. We now claim that X is also an important $A^* - B^*$ separator.

First, observe that X must be a minimal A^*-B^* separator: there exists a path from some vertex of A^* to each vertex $v \in X$ containing no other vertex of X, and similarly there exists a path from each vertex v of X to some vertex in $Y \subseteq B^*$ containing no other vertex of X. Now suppose that X is not an A^*-B^* important separator. Then there exists another A^*-B^* minimal separator X' with $|X'| \leq |X|$ with $R(X') \supset R(X)$. Since $X' \neq X$, there exists a vertex $v \in X \cap (V(G) \setminus X')$. Consider the set Y_v and the set of $|X|+1$ paths from X to Y_v which are vertex disjoint except at v, which appears twice. Since X' does not delete v, at least one of these paths from some vertex $x \in X$ to $t \in Y_v$ survives in $G \setminus X'$. By construction, there is a path P from some vertex $s \in A^*$ to $x \in X$, whose every internal vertex is not in X. Since every internal vertex of P is in $R(X)$, it must also be in $R(X')$. It follows that there is a path from s to x in $G \setminus X'$. Appending this to the path from x to t we obtain a path from s to t in $G \setminus X'$, contradicting the fact that X' is an A^*-B^* separator.

Finally, it is clear that the number of A^*-B^* important separators where $A^* \subseteq S$ and $B^* \subseteq T$ with $|A^*| \leq k$ and $|B^*| \leq 2k$ is at most 4^k times the number of ways of choosing A^* and B^*, and hence is at most $4^k \binom{|S|}{\leq k}\binom{|T|}{\leq 2k}$. Furthermore, once we fix A^* and B^*, all the important A^*-B^* separators of size $\leq k$ can be enumerated in time $\mathcal{O}(4^k \cdot k^2 \cdot (m+n))$, see Theorem 8.51 of [CFK+15]. The set of all subsets of A of size at most k and the set of all subsets of B of size at most k can be enumerated in time $\mathcal{O}(k\binom{|S|}{k})$ and $\mathcal{O}(k\binom{|T|}{k})$ respectively. This completes the proof. ☐

3 Detection Sets and Sample Sets in Directed Graphs

The goal of this section is to prove our result on detection sets and sample sets in directed graphs. The following theorems are our main results in this section.

Theorem 1.3. *For any directed graph G, given parameters $\epsilon \in (0,1)$ and $k \in \mathbb{N}$, there is an absolute constant c such that there is an (ϵ, k) detection set of size $f(k,\epsilon) = c\frac{k}{\epsilon}\log\frac{1}{\epsilon}$. Further, a random set of $f(k,\epsilon)$ vertices must be an (ϵ,k) detection set with probability at least $\frac{2}{3}$.*

Theorem 1.4. *For any directed graph G, given parameters $\epsilon \in (0,1)$ and $k \in \mathbb{N}$, there is an absolute constant c such that there is an (ϵ, k) sample set of size $f(k,\epsilon) = c\frac{k}{\epsilon^2}\log\frac{1}{\epsilon}$. Further, a random set of $f(k,\epsilon)$ vertices must be an (ϵ,k) sample with probability at least $\frac{2}{3}$.*

We start by defining a notion of (ϵ, k) nets and recalling the definitions for (ϵ, k) detection sets and (ϵ, k) sample sets from the overview. These definitions naturally extend those in [FM06] for undirected graphs.

Definition 3.1 (Net)**.** *Given a directed graph G, an (ϵ, k) net is a set of terminals $T \subseteq V(G)$ satisfying the following property: for every set of vertices F with $|F| \leq k$, the following two conditions are met:*

1. *For every SCC C of $G \setminus F$ with $|C| \geq \epsilon n$ we have*
$$|C \cap T| \geq 1.$$
2. *Let C^* be the union of all except one SCC in $G \setminus F$ such that $|C^*| \geq \epsilon n$, then we have*
$$|C^* \cap T| \geq 1.$$

Definition 3.2 (Detection Set)**.** *Given a directed graph G, an (ϵ, k) detection set is a set of terminals $T \subseteq V(G)$ satisfying the following property: for every set of vertices F with $|F| \leq k$ such that $V(G \setminus F)$ can be partitioned into (A, B) with $|A|, |B| \geq \epsilon n$ and there are no arcs from B to A, there exists $t_1, t_2 \in T$ such that there is no t_1-t_2 path in $G \setminus F$.*

Definition 3.3 (Sample Set)**.** *Given a directed graph G, an (ϵ, k) sample set is a set of terminals $T \subseteq V(G)$ satisfying the following property: for every set of vertices F with $|F| \leq k$, and every SCC C of $G \setminus F$, we have*
$$\left|\frac{|C \cap T|}{|T|} - \frac{|C|}{n}\right| \leq \epsilon.$$

We note that these definitions are almost identical to the one in [FM06], with the only difference being that we deal with directed graphs and strongly connected components in lieu of undirected graphs and connected components. Before proving the main results of this section, we show that (ϵ, k)-nets are also (ϵ, k)-detection sets. The proof is similar to that in undirected graphs.

Lemma 3.4. *Every (ϵ, k)-net is also an (ϵ, k) detection set.*

Proof. Let T be an (ϵ, k)-net for G. Consider a set of vertices F of size at most k, so that $V(G \setminus F)$ admits a partition A, B with $|A|, |B| \geq \epsilon n$ such that there are no arcs from B to A. Consider an SCC C of $G \setminus F$, and let C' be the union of all SCC's of $G \setminus F$ except C. Then clearly $|C'| \geq \epsilon n$, hence there exists a terminal $t_1 \in C' \cap T$. Let C_{t_1} be the SCC of $G \setminus F$ containing t_1. Now consider C'' which is the union of all SCC's of $G \setminus F$ except C_{t_1}. Again $|C''| \geq \epsilon n$, therefore there is a terminal $t_2 \in C'' \cap T$. But clearly t_1 and t_2 lie in different SCC's of $G \setminus F$, which means there is either no t_1-t_2 path or no t_2-t_1 path in $G \setminus F$, hence satisfying our detection set property. ∎

Thus we henceforth focus on obtaining (ϵ, k)-nets and samples.

First, we observe that showing a VC-dimension bound directly implies (ϵ, k) samples. This is similar to the approach in [FM06]. We start with a few definitions.

Definition 3.5 (Shattering a set of elements)**.** *Suppose we are given a set system $(\mathcal{S}, U)$ consisting of a family of sets $\mathcal{S}$, where each set in $\mathcal{S}$ consists of elements from a universe U. We say that a subset $W \subseteq U$ is* shattered *by $\mathcal{S}$, if for every subset $Y \subseteq W$, there exists a set $S \in \mathcal{S}$ so that $S \cap W = Y$.*

Definition 3.6 (VC-dimension)**.** *Given a set system $(\mathcal{S}, U)$, its VC-dimension is the size of the largest subset $W \subseteq U$ that can be shattered by $\mathcal{S}$.*

Theorem 3.7 (ϵ-net theorem, see [FM06])**.** *Let $(\mathcal{S}, U)$ be a set system with VC-dimension d, and universe size $|U| = n$. Then for every $\epsilon > 0$, there exists an absolute constant c such that a random subset $T \subseteq U$ of size $cd\frac{1}{\epsilon}\log\frac{1}{\epsilon}$ is an ϵ-net with probability at least $\frac{2}{3}$. Concretely, T satisfies $|S \cap T| \geq 1$ for every $S \in \mathcal{S}$ satisfying $|S| \geq \epsilon n$ with probability at least $\frac{2}{3}$.*

Theorem 3.8 (ϵ-sample theorem, see [FM06])**.** *Let $(\mathcal{S}, U)$ be a set system with VC-dimension d, and universe size $|U| = n$. Then for every $\epsilon > 0$, there exists an absolute constant c such that a random subset $T \subseteq U$ of size $cd\frac{1}{\epsilon^2}\log\frac{1}{\epsilon}$ is an ϵ-sample with probability at least $\frac{2}{3}$. Concretely, T satisfies $\left|\frac{|S|}{n} - \frac{|S\cap T|}{|T|}\right| \leq \epsilon$ for every $S \in \mathcal{S}$ with probability at least $\frac{2}{3}$.*

Let $\mathcal{S}$ be the family of sets consisting of all possible sets C which are either (a) a strongly connected component after deleting some vertex set F of size at most k or (b) the union of all but one strongly connected components after deleting some vertex set F of size at most k. Then it suffices to prove a VC-dimension bound of $\mathcal{O}(k)$ for this family of sets with the universe as the vertex set - this would then directly imply Theorems 1.3 and 1.4 using Theorems 3.7 and 3.8. Thus we will henceforth focus on upper-bounding the VC-dimension of the set system $\mathcal{S}$.

Towards this, we consider a slightly different problem of reachability from a single source after deleting k vertices.

Definition 3.9 (Single source reachability profile)**.** *Given a graph G, a source vertex s and a set of sink vertices T' with $s \notin T'$, we define the single source reachability profile of s as* $\text{Reach}^k(s, T') = \{R^G_{\{s\}}(F) \cap T' \mid F \subseteq V(G), |F| \leq k\}$ *In other words,* $\text{Reach}^k(s, T')$ *is the collection of subsets of T' reachable from s after deleting some set of at most k vertices from G.*

The next theorem bounds the size of $\text{Reach}^k(s, T')$.

Theorem 3.10. *Given a graph G, source vertex s and sink vertices T', $|\text{Reach}^k(s, T')| \leq 4^k\binom{|T'|}{\leq 2k}$.*

We prove Theorem 3.10 using our result on all-subset important separators, Theorem 1.1. We do this as follows. Given s and T', suppose there exists a set X of k vertices so that $P \subseteq T'$ is the subset of T' reachable from s in $G \setminus X$. Then in the following lemma, Lemma 3.11, we show that there is in fact an s-$(T' \setminus P)$ *important separator* X' of size $\leq k$, so that the subset of T' reachable from s in $G \setminus X$ remains P. The bound on all-subset important separators will then imply the bound on the reachability profile.

Lemma 3.11. *Suppose there exists a set X of $\leq k$ vertices so that $P \subseteq T'$ is the subset of T' reachable from s in $G \setminus X$. Then there is an s-$(T' \setminus P)$ important separator X' of size $\leq k$, so that the subset of T' reachable from s in $G \setminus X$ remains P.*

Proof. Without loss of generality, we assume that X is a minimal $s-(T' \setminus P)$ separator (otherwise we consider the subset of X that is minimal, this still has the same reachability set P). Let X' be the vertex mincut between X and $T' \setminus P$ that is closest to $T' \setminus P$. We have $|X'| \leq |X|$ since X itself is a vertex cut between X and $T' \setminus P$.

We show that the reachability profile after removing X' is also P. In other words, a vertex $t \in T'$ is reachable from s after removing X iff it is reachable after removing X'.

Consider some sink vertex $t \in T' \setminus P$, so that t is not reachable after removing X. Then, every s–t path contains a vertex in X. Since X' is a cut between X and $T' \setminus P$, this path must also contain a vertex in X'. This proves one direction.

Suppose that some vertex $t \in T'$ is not reachable after removing X'. Then, every s–t path contains a vertex in X'. Suppose for contradiction that some s–t path does not contain a vertex in X. Consider the prefix of the path leading to a vertex v in X'. Append to it a path from v to $T' \setminus P$ that only contains one vertex in X' (at the start) and no vertex of X. Such a path must exist since X' is a min-cut between X and $T' \setminus P$. The result is a path from s to $T \setminus P$ that does not contain a vertex in X. This contradicts the set P being the reachability profile after removing X.

Finally, we show that X' is an important s-$(T' \setminus P)$ separator. By Lemma 2.6 it is enough to show that X' is closest to $T' \setminus P$ and that X' is a minimal s-$(T' \setminus P)$ separator. It follows from the construction of X' that X' must be closest to $T' \setminus P$. Thus it is enough to show that X' is a minimal s-$(T' \setminus P)$ separator. Since X' is a min-cut between X and $T' \setminus P$, for every vertex $x' \in X'$, there exists a path Q from some vertex $x \in X$ to some vertex $t \in T' \setminus P$ through x' that does not contain any other vertex of X'. Since X is a minimal s-$(T' \setminus P)$ separator, there is a path P' from s to each vertex $x \in X$ that does not include any other vertex from X. Further P' does not contain any vertex from X': if it contains some vertex $y' \in X'$, it must be the case that $y' \notin X$. Then it follows that there is a path from s to y' in $G \setminus X$. Since X' is a min-cut between X and $T' \setminus P$, there is a path from each vertex of X' to some vertex of $T' \setminus P$ that does not have any internal vertex from X. In particular, there is a path from y' to some vertex of $T' \setminus P$ which does not include any vertex of X. This path, when appended to the path from s to y' in $G \setminus X$, results in a path from s to some vertex of $T' \setminus P$ which does not contain any vertex of X, which in turn contradicts the fact that X is a cut between s and $T' \setminus P$. Finally, appending P' with Q gives a path from s to $t \in T' \setminus P$ that contains x', while containing no other vertex of X'. Since such a path exists for any $x' \in X'$, it follows that X' is minimal, and hence X' is an important s-$T' \setminus P$ separator. □

Proof of Theorem 3.10. Suppose $P \in \text{Reach}^k(s, T')$. Lemma 3.11 implies that there is an important s-$(T' \setminus P)$ separator X so that the reachability set of s in $G \setminus X$ is exactly $P \subseteq T'$. Using Theorem 1.1 the number of such important separators is at most $4^k \binom{|T'|}{\leq 2k}$, and hence the number of such P is at most $4^k \binom{|T'|}{\leq 2k}$. □

Next, we show why Theorem 3.10 implies a VC-dimension bound.

Theorem 3.12. *Let $\mathcal{S}$ be the family of sets over consisting of all possible sets C which are either (a) a strongly connected component after deleting some vertex set F of size at most k or (b) the union of all but one strongly connected components after deleting some vertex set F of size at most k. Then the VC-dimension of $\mathcal{S}$ is $\mathcal{O}(k)$.*

Proof. From the definition of VC-dimension, it is enough to show that there no set of size $\geq dk$ that can be shattered, for some constant d. Suppose a set of terminals T can be shattered. We will show that $|T| = \mathcal{O}(k)$.

Consider some $P \subseteq T$, with $P \neq \emptyset$, and suppose that after deleting some subset F of vertices of size at most k, some strongly connected component C of $G \setminus F$ satisfies $C \cap T = P$. Fix any $s \in C \cap T$, and let $T' = T \setminus \{s\}$. Let G^R denote the graph G with arcs reversed. Then by definition of an SCC, $P \setminus s$ is the exact subset of T' reachable from s in both $G \setminus F$ and $G^R \setminus F$. But this means $P \setminus s = P_1 \cap P_2$ for some $P_1 \in \mathrm{Reach}^k_G(s, T')$ and $P_2 \in \mathrm{Reach}^k_{G^R}(s, T')$.

Together, we obtain that if $P \subseteq T$ is $C \cap T$ for some SCC C after deleting some k vertices, there must exist a vertex $s \in P$ such that $P \setminus \{s\} = P_1 \cap P_2$ for some $P_1 \in \mathrm{Reach}^k_G(s, T \setminus \{s\})$ and $P_2 \in \mathrm{Reach}^k_{G^R}(s, T \setminus \{s\})$. But then the number of choices for P is at most $|T|(4^k \binom{|T|}{\leq 2k})^2$.

Alternatively, suppose that $P = C' \cap T$ where C' is the union of all but one strongly connected component, say C^*, after deleting some set F of k vertices. Suppose $P \neq T$, so that $C^* \cap T \neq \emptyset$. Consider some $s' \in C^* \cap T$. Then it follows that $(T \setminus P) \setminus \{s'\} = Q_1 \cap Q_2$ where $Q_1 \in \mathrm{Reach}^k_G(s', T \setminus \{s'\})$ and $Q_2 \in \mathrm{Reach}^k_{G^R}(s', T \setminus \{s'\})$. But then the number of choices for $T \setminus P$ is at most $|T|(4^k \binom{|T|}{\leq 2k})^2$, and hence the number of choices for P is at most $|T|(4^k \binom{|T|}{\leq 2k})^2$.

Thus the total number of choices for P across both cases is at most $2|T|(4^k \binom{|T|}{\leq 2k})^2 + 2$ (accounting for $P = \emptyset$ in the first case, and $P = T$ in the second) which is not less than $2^{|T|}$ only when $|T| = \mathcal{O}(k)$, finishing the proof. □

4 FPT Algorithm for Finding Directed Balanced Separators

The goal of this section is to prove our result on Directed Balanced Separator.

Theorem 1.6. *Given an instance of* Directed Balanced Separator *there is a randomized algorithm, that for any $\epsilon > 0$, runs in time $2^{\mathcal{O}(k \min\{\log \frac{1}{b}, \log k\} \log \frac{1}{\epsilon}/\epsilon^2)}(m+n)$ and with constant probability outputs either (a) a set of vertices (arcs) F' of size at most k such that every strongly connected component of $G \setminus F'$ has size at most $(b+\epsilon)n$ or (b) concludes correctly that there is no b-balanced separator of size at most k.*

We will crucially exploit our results on sample sets, Theorem 1.4. We will also need the following result, which is about computing skew separators. The Skew Separator problem is a special case of Directed Multicut.

Definition 4.1 (Skew Separator [CLL+08])**.** *Given a directed graph G and a set of ℓ terminal pairs $\{(s_i, t_i)\}_{i \in [\ell]}$ and an integer k, find a set of vertices F of size at most k, so that $G \setminus F$ has no $s_i - t_j$ path for $i \geq j$, $i, j \in [\ell]$.*[5]

The analagous edge version, Skew Edge Separator can be defined similarly.

Theorem 4.2 (Extension of Theorem 8.41 of [CFK+15])**.** Skew Separator *and* Skew Edge Separator *admit FPT algorithms running in time $\mathcal{O}(4^k \cdot k^3 \cdot (n+m))$.*

Proof. Theorem 8.41 of [CFK+15] is for the edge version Skew Edge Separator. For the sake of completeness, we reduce the vertex version to the edge version. The reduction is standard and as follows.

Given an instance of Skew Separator, we create an instance of Skew Edge Separator. Define graph G' as follows. For each vertex $v \in V(G)$, create two vertices v_{in} and v_{out} in $V(G')$. For every edge $\{u, v\} \in E(G)$, add an edge between u_{out} and v_{in} in $E(G')$. Finally, add an edge u_{in} to u_{out} to $E(G')$ for each $u \in V(G)$. The instance of Skew Edge Separator is then the graph G', the integer k and the set of pairs $(s_{i_{in}}, t_{i_{out}})$ for each $i \in [\ell]$.

[5] Again, we note that we can delete terminals as well to form the solution.

We claim that there for any integer k', there is a solution to Skew Separator of size at most k' in G iff there is a solution to Skew Edge Separator of size at most k' in G'. First we prove the forward direction. Let X be a set of vertices that forms a solution to the Skew Separator instance. Then clearly by our construction, the set of edges (x_{in}, x_{out}) for $x \in X$ forms a solution to the Skew Edge Separator instance.

For the backward direction, let X' be a set of edges that forms a solution in G'. If an edge in X' is of the form (y_{out}, x_{in}), we replace it by the edge (y_{out}, y_{in}), so that now every edge in X' is of the form (y_{out}, y_{in}). Then let $Y = \{y \mid (y_{in}, y_{out}) \in X'\}$. Clearly, $|Y| \leq k$, and Y must form a solution in G for the Skew Separator instance. □

Lemma 4.3. *Let T be an (ϵ, k) sample set. Suppose that the* Directed Balanced Separator *instance is a YES-instance, so that there exists a set of vertices (arcs) F with $|F| \leq k$ whose deletion leaves every strongly connected component with at most bn vertices. Then one can in $\mathcal{O}(2^{\mathcal{O}(|T| \min\{\log \frac{1}{b}, \log |T|\})} m)$ time, find a set of vertices (arcs) F' with $|F'| \leq k$ so that there is a partition (X, Y) of $G \setminus F'$ such that every SCC of $G \setminus F'$ has size at most $(b + \mathcal{O}(\epsilon))n$.*

Proof. Broadly, our goal is to obtain a set of vertices (arcs) F' with $|F'| \leq k$, so that in $G \setminus F'$, every strongly connected component does not have more than $(b+\epsilon)|T|$ terminals. The sample set property would then imply that every strongly connected component of $G \setminus F'$ would have size at most $(b + 2\epsilon)n$.

We proceed as follows. Let C be a strongly connected component of $G \setminus F$. Then by the sample set property, we must have $|\frac{n}{|T|}|C \cap T| - |C|| \leq \epsilon n$, which in turn means that $|C \cap T| \leq (b + \epsilon)|T|$. Consider the strongly connected components of $G \setminus F$. They form a DAG. We sort this DAG in the topological order - let $(C_1, C_2 \ldots C_l)$ be the ordering of strongly connected components such that there is no arc from C_j to C_i whenever $j > i$, for $i, j \in [\ell]$.

- **Case 1:** We assume that $\log \frac{1}{b} \leq \log |T|$. Partition the indices $\{1, 2 \ldots \ell\}$ into $t \leq \frac{1}{b}$ intervals $[l_i, r_i]_{i=1}^{t}$ as follows. Let $l_1 := 1$. For each i, r_i is defined as the smallest integer greater than or equal to l_i for which $\bigcup_{j=l_i}^{r_i} C_j$ contains at least $b|T|$ terminals. If no such integer exists, then $r_i := \ell$ and we set $t := i$. If $r_i \neq \ell$ we set $l_{i+1} := r_i + 1$ and continue.

  For each terminal $t \in T$, we guess if $t \in F$, and if not, we guess i such that $t \in \bigcup_{j=l_i}^{r_i} C_j$. We also guess if $t \in C_{r_i}$. All these guesses together take time $2^{\mathcal{O}(|T| \log \frac{1}{b})}$.

  We now create a Skew Separator (Skew Edge Separator) instance as follows. Let $T_{i1} = ((\bigcup_{j=l_i}^{r_i} C_j) \setminus C_{r_i}) \cap T$ and $T_{i2} = C_{r_i} \cap T$ for each $i \in [\ell]$. Observe that by construction, for any $i \in [\ell]$, both T_{i1} and T_{i2} have at most $(b+\epsilon)|T|$ terminals. Next, contract T_{i1} into a single vertex t_{i1}, and contract T_{i2} into a single vertex t_{i2}. Let G^C denote the contracted graph. The instance of Skew Separator is the graph G^C and the pairs $\{(t_{11}, t_{12}), (t_{12}, t_{21}), (t_{21}, t_{22}) \ldots (t_{\ell 1}, t_{\ell 2})\}$. (If some $T_{i1} = \emptyset$, we replace the pairs $(t_{i-1,2}, t_{i1}), (t_{i1}, t_{i2})$ by the single pair $(t_{i-1,2}, t_{i2})$, we omit this detail.)

  Skew Separator (Skew Edge Separator) then returns a set of vertices (arcs) F' with $|F'| \leq k$, so that in $G \setminus F'$, there is no path from $t_{c_1 d_1} \to t_{c_2 d_2}$ whenever $c_2 \leq c_1$ or when $c_2 = c_1$ and $d_2 \leq d_1$, for $c_1, c_2 \in [\ell]$ and $d_1, d_2 \in \{1, 2\}$. In particular, this means that no strongly connected component of $G \setminus F'$ has more than $(b + \epsilon)|T|$ terminals, since for any $c \in [\ell]$ and $d \in \{1, 2\}$, $|T_{cd}| \leq (b + \epsilon)|T|$.

- **Case 2:** We now consider the case $\log \frac{1}{b} \geq \log |T|$, so that $|T| < \frac{1}{b}$. Observe that in the above analysis, the number of indices i satisfying $C_i \cap T \neq 0$ is at most $|T|$. We guess the ordered partition $(W_1, W_2 \ldots W_p)$ of $T \setminus F'$, where $p \leq |T|$, such that for each $j \in [p]$, $W_j = C_i \cap T$

where C_i is the j^{th} component in $(C_1, C_2 \dots C_\ell)$ satisfying $C_i \cap T \neq \emptyset$. In other words, we guess the ordered partition of the terminal set induced by the strongly connected components $(C_1, C_2 \dots C_\ell)$ of $G \setminus F$. This can be accomplished in time $|T|^{|T|} = 2^{|T| \log |T|}$.

We now create a Skew Separator (Skew Edge Separator) instance as follows. For each $j \in [p]$, contract W_j into a single vertex w_j. Let H^C denote the contracted graph. The instance of Skew Separator is the graph H^C and the pairs $\{(w_1, w_2), (w_2, w_3), (w_3, w_4) \dots (w_{p-1}, w_p)\}$.

Skew Separator (Skew Edge Separator) then returns a set of vertices (arcs) F' with $|F'| \leq k$, so that in $G \setminus F'$, there is no path from $w_{c_1} \to w_{c_2}$ whenever $c_2 \leq c_1$. In particular, this means that no strongly connected component of $G \setminus F'$ has more than $(b+\epsilon)|T|$ terminals, since for any $c \in [p]$, $|W_c| \leq (b+\epsilon)|T|$.

Finally, by the property of (ϵ, k) sample sets (Definition 3.3), in either case, it follows that since any strongly connected component of $G \setminus F'$ has at most $(b+\epsilon)|T|$ terminals, it must have at most $(b+2\epsilon)n$ vertices. Note that though the statement of Definition 3.3 only deals with the case when F' is a set of vertices, we have the simple observation that if F' is a set of arcs and C is a strongly connected component after deleting F', it is also a strongly connected component after deleting the endpoints of these arcs not in C. This completes the proof.

□

The proof of Theorem 1.6 now follows immediately from Theorem 1.4 and Lemma 4.3.

5 Approximation Algorithm for Directed Balanced Separators

In this section, the goal is to prove the following result, which shows a $\mathcal{O}(\sqrt{\log k})$ approximation for Directed Balanced Separator.

Theorem 5.1. *There is an $\mathcal{O}(\sqrt{\log k})$ approximation to* Directed Balanced Separator *in polynomial time. Formally, given an instance of* Directed Balanced Separator *with $b = \Omega(1)$, suppose there is a set of vertices (arcs) F with $|F| \leq k$, so that every strongly connected component of $G \setminus F$ has at most bn vertices. Then there is a polynomial time randomized algorithm that with constant probability finds a set of vertices (arcs) F' with $|F'| \leq \mathcal{O}(k\sqrt{\log k})$ so that in $G \setminus F'$, every strongly connected component has size at most $b'n$ for some $b' < 1$ depending on b.*

Our result is obtained by essentially following the algorithm of [ACMM05] together with our theorem on sample sets, Theorem 3.8. On a high level, the reason we get $\mathcal{O}(\sqrt{\log k})$ approximation comes down to the fact that one can find a balanced separator with respect to a sample set T, which will automatically be a balanced separator for the entire graph as well due to the property guaranteed by sample sets. Once we have this equivalence, in order to obtain a balanced separator with respect to the set T, we now need to use the ARV structure theorem [ARV09] only on the sample set vectors, which are at most $|T|$ in number. By Theorem 1.4 there is such a sample set T of size at most $\mathcal{O}(k)$. This helps replace the $\mathcal{O}(\sqrt{\log n})$ factor by $\mathcal{O}(\sqrt{\log k})$. However, for the sake of completeness and clarity we give the full algorithm by mostly following the algorithm of [ACMM05] with a few modifications/simplifications.

We will restrict our attention to edge cuts, as again the standard reduction in Theorem 4.2 can easily reduce the vertex version to the edge version.

Notice that if a separator F is b-balanced for some $b = \Omega(1)$, then every strongly connected component of $G \setminus F$ has size at most bn. This in turn means that using a prefix of the topological order of the strongly connected components in $G \setminus F$, we can obtain sets A, B with $|A|, |B| \geq c'n$

for $c' = \frac{1-b}{2}$, such that there are no arcs from A to B in $G \setminus F$. Thus, upto constant factors in the balance, one can equivalently think of the Directed Balanced Separator problem as that of finding a set of vertices (arcs) F with $|F| \leq k$ and a partition A, B of $G \setminus F$, so that $|A|, |B| \geq c'n$ and there are no arcs from A to B in $G \setminus F$.

Now we consider the terminal version of the problem, to which we will reduce via sample sets later. Given a terminal set T, our revised goal is to solve the following problem: given a parameter $c = \Omega(1)$, find a set of edges F with $|F| \leq k$ so that there is a partition of $G \setminus F$ into A, B such that $|A \cap T|, |B \cap T| \geq c|T|$ and there are no arcs from A to B in $G \setminus F$. In what follows, we will mostly follow the algorithm of [ACMM05] with small modifications which we state as and when required.

We assume that the vertices are labelled $\{1, 2, \ldots n\}$, and add a special vertex 0, which will be a non-terminal and will be the reference vertex for the "A- side" of the cut. We start with the SDP relaxation (extended in our case to the terminal version) in [ACMM05] and follow their algorithm.

$$\min \frac{1}{8} \sum_{e=\{i,j\} \in E(G)} d(i,j)$$

$$\|v_i\|^2 = 1 \quad \forall i \in [n] \cup \{0\}$$

$$\|v_i - v_j\|^2 + \|v_j - v_k\|^2 \geq \|v_i - v_k\|^2 \quad \forall i, j, k \in [n] \cup \{0\}$$

$$\sum_{i<j, i,j \in T} \|v_i - v_j\|^2 \geq 4c(1-c)|T|^2$$

Here the "directed distance" $d(i,j) = \|v_i - v_j\|^2 + \|v_j - v_0\|^2 - \|v_i - v_0\|^2$ is as defined in [ACMM05]. The cannonical solution for this SDP is obtained by setting $v_i = v_0$ for each $i \in A$ and $v_i = -v_0$ for each $i \in B$.

Theorem 5.2 (ARV separation theorem [ARV09], as stated in [ACMM05])**.** *Given a set of ℓ_2^2 unit vectors v_i, $i \in [n]$, such that $\sum_{i<j} \|v_i - v_j\|^2 \geq 4c(1-c)n^2$, there exists a polynomial time algorithm that finds disjoint sets L, R with $|L|, |R| = \Omega(n)$ such that for any $i \in L, j \in R$, we have $\|v_i - v_j\|^2 \geq \Omega(1/\sqrt{\log n})$.*

Now we follow Algorithm 4 of [ACMM05]. Step 1 solves the SDP. In Step 2, given the vectors from the SDP, we apply the ARV separation theorem to only the vectors v_i, $i \in T$ to find $\Delta = \frac{1}{\sqrt{\log |T|}}$ separated sets L and R. Concretely, we obtain disjoint sets $L, R \subseteq T$ with $|L|, |R| = \Omega(|T|)$ so that $\|v_i - v_j\|^2 \geq \Delta$ for any $i \in L$ and $j \in R$.

Next, we define r so that both $L^+ = \{i \in L \mid |v_0 - v_i|^2 \leq r^2\}$ and $L^- = \{i \in L \mid |v_0 - v_i|^2 \geq r^2\}$ have more than $\frac{|L|}{2}$ vertices. Note that such an r always exists. Once we fix r, we define $R^+ = \{i \in R \mid |v_0 - v_i|^2 \leq r^2\}$ and $R^- = \{i \in R \mid |v_0 - v_i|^2 \geq r^2\}$ similarly. Finally, if $|R^+| > \frac{|R|}{2}$ we compute a directed min-cut F between R^+ and L^-, else we compute a directed min-cut F between L^+ and $R-$.

Lemma 5.3. *$|F| \leq \mathcal{O}(\frac{1}{\Delta} SDP)$ and F is a b^*-balanced separator with respect to T for some $b^* < 1$ depending on c, where SDP is the optimal SDP value.*

Proof. Without loss of generality, we assume that $|R^+| > \frac{|R|}{2}$. The other case is symmetric. First, observe that for any $i \in R^+$, $j \in L^-$, we have $d(i,j) > \Delta$. This follows from the definition of the sets R^+ and L^- which in particular means that $|v_i - v_0|^2 \leq r^2$ and $|v_j - v_0|^2 \geq r^2$ and hence $d(i,j) = \|v_i - v_j\|^2 + \|v_j - v_0\|^2 - \|v_i - v_0\|^2 \geq \|v_i - v_j\|^2 \geq \Delta$.

To show the bound on number of cut edges, we analyze the algorithm as follows. For sets U and W, we define $d(U,W) := \min_{u\in U, w\in W} d(u,w)$. Similarly, one can define $d(U,w) := d(U,\{w\})$ and $d(u,W) = d(\{u\},W)$ for $u,w \in [n]\cup\{0\}$. Pick a random threshold $\tau \in (0, \frac{\Delta}{4})$ at random and consider the set $X = \{x \mid d(R^+,x) \le \tau\}$. Clearly X includes R^+, and does not include L^-. Also for each arc (a,b) such that $a \in X$ and $b \notin X$, we have $d(a,b) \ge d(R^+,b) - d(R^+,a)$ since the directed distance metric d satisfies the triangle inequality. It follows that the probability that the edge (a,b) is cut is at most $\mathcal{O}(\frac{d(a,b)}{\Delta})$ and hence the number of edges cut in expectation is at most $\mathcal{O}(\frac{SDP}{\Delta})$. In particular, the (R^+,L^-) min-cut has size at most $\mathcal{O}(\frac{SDP}{\Delta})$.

Next, we show that this cut is b^*- balanced with respect to T, for some $b^* < 1$ depending only on b. This follows directly from the description of the algorithm. Recall that the ARV separation theorem ensures that $|L|, |R| \ge \Omega_c(|T|)$. Since $|R^+| > \frac{|R|}{2}$ and by construction $|L^-| > \frac{|L|}{2}$ and there is no path from R^+ to L^- in $G \setminus F$, the balance condition follows. □

Proof of Theorem 5.1. Given the directed graph G which forms the instance of Directed Balanced Separator, we first guess the value of k, so that there is a set of edges F with $|F| \le k$ so that each strongly connected component of $G \setminus F$ has size at most bn. Using a prefix of the strongly connected components of $G \setminus F$, it follows that there is a partition $X \cup Y$ of $V(G)$ so that $|X|, |Y| \ge c'n$ where $c' = \frac{1-b}{2}$ and there is no arc from X to Y in $G \setminus F$. Next, construct a (ϵ, k) sample set T for some constant $\epsilon \ll c'$. With constant probability, T is a sample set and $|T \cap X|, |T \cap Y| \ge (c' - \mathcal{O}(\epsilon))|T| \ge \frac{c'}{2}|T|$. Use the above algorithm and Lemma 5.3 with $c = \frac{c'}{2}$ to obtain a set of edges F' with $|F'| \le \mathcal{O}(k\sqrt{\log |T|})$ that forms a balanced separator for T. It is clear that in $G \setminus F'$, every strongly connected component has at most $b^*|T|$ terminals where b^* depends only on c', as given by Lemma 5.3. We choose ϵ small enough so that $b^* < 1 - 2\epsilon$. It follows from Definition 3.3 that each strongly connected component has size at most $(b^* + \epsilon)n$. Finally, note that $|T| = \mathcal{O}(k)$ (since ϵ is a constant), hence $\sqrt{\log |T|} = \mathcal{O}(\sqrt{\log k})$. This concludes the proof.

□

6 Vertex Cut Sparsifiers

In this section, we present our result on vertex cut sparsifiers. We start with a few definitions.

Definition 6.1. *Given a directed graph G, a set of terminals $T \subseteq V(G)$, and an integer c, a directed graph G' satisfying $T \subseteq V(G')$ is said to be a (c,T) vertex cut sparsifier for G if for any partition $A \cup B$ of T such that the size of the minimum A-B (vertex) cut in G is at most c, the size of the A-B mincut is the same in G and G'.*

Our next result gives our deterministic construction of vertex sparsifiers in directed graphs.

Theorem 1.9. *Given a directed graph G, a set of terminals T and integer c, there is a deterministic algorithm that runs in time $\mathcal{O}(\psi(|T|,c)\cdot(m+n))$ and computes a (c,T) vertex sparsifier G' for G of size at most $\psi(|T|,c)$, where $\psi(|T|,c) = \binom{|T|}{\le 3c}2^{\mathcal{O}(c)}$. Additionally, G' satisfies $T \subseteq V(G') \subseteq V(G)$, with the property that for every partition $A \cup B$ of T and every subset $X \subseteq V(G)$ with $|X| \le c$, X is an A-B important separator in G if and only if $X \subseteq V(G')$ and X is an A-B important separator in G'.*

To prove this result, we use our result on important separators combined with the standard approach of "closing" unnecessary vertices - given a graph G and a vertex $v \in V(G)$, applying the closure operation to v gives the graph G' with the vertex set $V(G) \setminus \{v\}$. The arc set of G' is the

same as that of G, except that it excludes all arcs incident on v, and for every pair of vertices (u, w) such that u is an in-neighbour of v and w is an out-neighbour of v, we add the arc (u, w) to $E(G')$. Before we prove this result, we have the following basic observation. Recall that given a graph H and sets $S, Z \subseteq V(H)$, $R_H^S(Z)$ denotes the set of vertices reachable from some vertex of S in $H \setminus Z$.

Lemma 6.2. *Let $A, Y \subseteq V(G)$ and let G' be the graph obtained by applying the closure operation on a vertex $v \in V(G)$. Also suppose $v \notin Y$. Then,*

1. $R_{G'}^A(Y) \subseteq R_G^A(Y)$.
2. $R_G^A(Y) \setminus \{v\} \subseteq R_{G'}^A(Y)$

Proof. To prove the first part, consider a $y \in R_{G'}^A(Y)$. Consider the path from some vertex $a \in A$ to y in $G' \setminus Y$. If this path uses an arc (z_1, z_2) of G' that is not an arc of G, then we use v to replace this arc by the 3-path $z_1 v z_2$ and obtain a new path from a to y in $G \setminus Y$.

For the second part, similarly consider a $y \in R_G^A(Y) \setminus \{v\}$. Consider the path from some $a \in A$ to y in $G \setminus Y$. If this path uses some arc that goes through v, consider the predecessor z_1 and successor z_2 of v in this path: there must be an arc between z_1 and z_2 by the closure operation. We replace the 3-path $z_1 v z_2$ by the arc (z_1, z_2) to obtain a path from a to y in $G' \setminus Y$. □

The next two lemmas show that applying the closure operation to a single vertex which is not in any A-B important separator of size at most c, across all partitions $A \cup B$ of T, preserves the set of important separators. This in turn will imply that we can apply the closure operation at once to all such vertices.

Lemma 6.3. *Given a directed graph G, disjoint sets A, B and a vertex $v \notin A \cup B$ which is not in any A-B important separator of size $\leq c$, let G' be the graph obtained by closing v. Then if $X \subseteq V(G')$, $|X| \leq c$, is an important A-B separator in G', then it is an A-B important separator in G as well.*

Proof. We need to prove three things. First, that X is an A-B separator in G. Second, it is inclusion-wise minimal. Third, that X is not *dominated* by any other separator - that is, there does not exist another separator Y with $|Y| \leq |X|$ so that $R_G^A(X) \subset R_G^A(Y)$.

First, X is a separator in G, for if not, an a-b path for dome $a \in A$ and $b \in B$ in $G \setminus X$ would imply such a path in $G' \setminus X$ by Lemma 6.2, a contradiction.

Next, if $X \setminus \{w\}$ for some $w \in X$ is also an A-B separator in G, then it must be an A-B separator in G' as well: for if there is a a-b path in $G' \setminus (X \setminus \{w\})$ for some $w \in X, a \in A, b \in B$, then there must be an a-b path in $G \setminus (X \setminus \{w\})$ as well by Lemma 6.2.

Finally, suppose Y is an important separator that dominates X in G. Observe that $v \notin Y$. First we show that Y is a separator in G'. This follows since an a-b path in $G' \setminus Y$ implies an a-b path in G by Lemma 6.2.

Now we wish to prove that Y dominates X in G' as well. Consider a vertex $u \in R_A^{G'}(X)$. Then by Lemma 6.2 it is clear that $u \in R_A^G(X)$ as well. But this means $u \in R_A^G(Y)$, by our assumption. This in turn means that $u \in R_A^{G'}(Y)$. Thus $R_A^{G'}(X) \subseteq R_A^{G'}(Y)$. Finally, we note that X was an inclusion-wise minimal separator in G', and $Y \neq X$, therefore there must exist at least one vertex of X that is reachable in $R_A^{G'}(Y)$. This means that in fact $R_A^{G'}(X) \subset R_A^{G'}(Y)$, and we obtain that Y dominates X in G', a contradiction to the fact that X was important in G'. □

The next lemma is analogous and proves the other direction.

Lemma 6.4. *Given a directed graph G, disjoint sets A, B and a vertex v which is not in any $A-B$ important separator of size $\leq c$, let G' be the graph obtained by closing v. Then if $X \subseteq V(G)$, $|X| \leq c$ is an A-B important separator in G, then it is an A-B important separator in G' as well.*

Proof. We need to prove three things. First, that X is an A-B separator in G'. Second, it is inclusion-wise minimal. Third, that X is not dominated by any other separator - that is, there does not exist another separator Y with $|Y| \leq |X|$ so that $R^A_{G'}(X) \subset R^A_{G'}(Y)$.

First, we show that X is an A-B separator in G'. If not, there is an a-b path in $G' \setminus X$ for some $a \in A, b \in B$. But then there is an a-b path in $G \setminus X$ by Lemma 6.2, a contradiction.

Second, suppose $X \setminus \{w\}$ is an A-B separator in G' for some $w \in X$. Then $X \setminus \{w\}$ is also an A-B separator in G, for if there is an a-b path in $G \setminus (X \setminus \{w\})$, then such a path exists in $G' \setminus (X \setminus \{w\})$ as well by Lemma 6.2, again a contradiction.

Third, suppose that X is dominated by Y in G'. First, we show that Y is a separator in G. This holds since any a-b path in $G \setminus Y$ would imply an a-b path in $G' \setminus Y$ by Lemma 6.2.

Then we would like to show that X is dominated by Y in G as well. Pick some $u \in R_G(X)$. If $u \neq v$, then it is clear that $u \in R_{G'}(X)$ as well by Lemma 6.2. Then $u \in R_{G'}(Y)$ by our assumption, and it follows that $u \in R_G(Y)$ by Lemma 6.2. Now suppose $v \in R_G(X)$. Pick any in-neighbour v^* of v in G, so that v^* in $R_G(X)$. Then it follows that $v^* \in R_G(Y)$ by our previous argument. But v is an out-neighbour of v^*, and $v \notin Y$. Hence it follows that $v \in R_G(Y)$. Thus it follows that $R_G(X) \subseteq R_G(Y)$. Also since X was an inclusion-wise minimal separator in G and $Y \neq X$, at least one vertex of X must be reachable in $G \setminus Y$, which means that in fact $R_G(X) \subset R_G(Y)$ which contradicts the fact that X is an important separator in G. □

Proof of Theorem 1.9. Given T, consider all pairs of subsets (A', B') of T satisfying $|A'| \leq k$ and $|B'| \leq 2k$. Let Y be the union of all (A', B') important separators in G across all such (A', B'). Observe that $|Y| \leq 4^k \binom{|T|}{\leq 3k} \binom{3k}{\leq k}$.

Apply the closure operation to every vertex of $Z = V(G) \setminus Y$ simultaneously. More precisely, define $Z_{in} = \{v \in V(G) \mid (v, z) \in E(G) \text{ for some } z \in Z\}$ and $Z_{out} = \{v \in V(G) \mid (z, v) \in E(G) \text{ for some } z \in Z\}$. Then in G, we delete the set of vertices Z, and add arcs (z_{in}, z_{out}), $z_{in} \in Z_{in}, z_{out} \in Z_{out}$ to obtain G'.

The correctness and running time of the algorithm follows from the soundness of Theorem 1.1, Lemma 6.3 and Lemma 6.4. □

References

[ACMM05] Amit Agarwal, Moses Charikar, Konstantin Makarychev, and Yury Makarychev. $\mathcal{O}(\sqrt{\log n})$ approximation algorithms for min uncut, min 2cnf deletion, and directed cut problems. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 573–581, 2005. 4, 5, 10, 18, 19

[ARV09] Sanjeev Arora, Satish Rao, and Umesh Vazirani. Expander flows, geometric embeddings and graph partitioning. *Journal of the ACM (JACM)*, 56(2):1–37, 2009. 4, 5, 10, 18, 19

[CDK+21] Parinya Chalermsook, Syamantak Das, Yunbum Kook, Bundit Laekhanukit, Yang P Liu, Richard Peng, Mark Sellke, and Daniel Vaz. Vertex sparsification for edge connectivity. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1206–1225. SIAM, 2021. 6

[CFK+15] Marek Cygan, Fedor V Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized algorithms*, volume 4. Springer, 2015. 1, 11, 12, 16

[CHM13] Rajesh Chitnis, MohammadTaghi Hajiaghayi, and Dániel Marx. Fixed-parameter tractability of directed multiway cut parameterized by the size of the cutset. *SIAM Journal on Computing*, 42(4):1674–1696, 2013. 1, 2

[Chu12] Julia Chuzhoy. On vertex sparsifiers with steiner nodes. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 673–688, 2012. 6

[CKL+20] Marek Cygan, Paweł Komosa, Daniel Lokshtanov, Marcin Pilipczuk, Michał Pilipczuk, Saket Saurabh, and Magnus Wahlström. Randomized contractions meet lean decompositions. *ACM Transactions on Algorithms (TALG)*, 17(1):1–30, 2020. 1, 4

[CLL+08] J Chen, Y Liu, S Lu, B O'Sullivan, and Igor Razgon. A fixed-parameter algorithm for the directed feedback vertex set problem. *Journal of the ACM*, 55(5):1–19, 2008. 1, 9, 16

[CLL09] Jianer Chen, Yang Liu, and Songjian Lu. An improved parameterized algorithm for the minimum node multiway cut problem. *Algorithmica*, 55(1):1–13, 2009. 1, 2

[CLP+19] Marek Cygan, Daniel Lokshtanov, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. Minimum bisection is fixed-parameter tractable. *SIAM Journal on Computing*, 48(2):417–450, 2019. 1, 2, 4

[DF+13] Rodney G Downey, Michael R Fellows, et al. *Fundamentals of parameterized complexity*, volume 4. Springer, 2013. 1

[ENRS99] Guy Even, Joseph Naor, Satish Rao, and Baruch Schieber. Fast approximate graph partitioning algorithms. *SIAM Journal on Computing*, 28(6):2187–2214, 1999. 4, 5

[FHL08] Uriel Feige, Mohammadtaghi Hajiaghayi, and James R Lee. Improved approximation algorithms for minimum weight vertex separators. *SIAM Journal on Computing*, 38(2):629–657, 2008. 4

[FM06] Uriel Feige and Mohammad Mahdian. Finding small balanced separators. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 375–384, 2006. 1, 2, 3, 4, 5, 8, 9, 13, 14

[FY03] Uriel Feige and Orly Yahalom. On the complexity of finding balanced oneway cuts. *Information processing letters*, 87(1):1–5, 2003. 4

[HLW21] Zhiyang He, Jason Li, and Magnus Wahlström. Near-linear-time, optimal vertex cut sparsifiers in directed acyclic graphs. In *29th Annual European Symposium on Algorithms (ESA 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021. 6, 10

[Kle03] Jon Kleinberg. Detecting a Network Failure. *Internet Mathematics*, 1(1):37 – 55, 2003. 2

[KSS08] Jon Kleinberg, Mark Sandler, and Aleksandrs Slivkins. Network failure detection and graph connectivity. *SIAM Journal on Computing*, 38(4):1330–1346, 2008. 2

[KW12] Stefan Kratsch and Magnus Wahlström. Representative sets and irrelevant vertices: New tools for kernelization. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 450–459. IEEE, 2012. 6, 10

[KW20] Stefan Kratsch and Magnus Wahlström. Representative sets and irrelevant vertices: New tools for kernelization. *Journal of the ACM (JACM)*, 67(3):1–50, 2020. 6

[Liu23] Yang P Liu. Vertex sparsification for edge connectivity in polynomial time. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, volume 251, page 83. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. 6

[LM13] Daniel Lokshtanov and Dániel Marx. Clustering with local restrictions. *Information and Computation*, 222:278–292, 2013. 1

[LMR+21] Daniel Lokshtanov, Pranabendu Misra, MS Ramanujan, Saket Saurabh, and Meirav Zehavi. Fpt-approximation for fpt problems. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 199–218. SIAM, 2021. 1, 2

[LN22] Jason Li and Jesper Nederlof. Detecting feedback vertex sets of size k in $\mathcal{O}(2.7^k)$ time. *ACM Transactions on Algorithms (TALG)*, 18(4):1–26, 2022. 1

[LRS+25] Daniel Lokshtanov, Maadapuzhi-Sridharan Ramanujan, Saket Saurabh, Roohani Sharma, and Meirav Zehavi. Wannabe bounded treewidth graphs admit a polynomial kernel for directed feedback vertex set. *ACM Trans. Comput. Theory*, 17(1), February 2025. 8

[Mar06] Dániel Marx. Parameterized graph separation problems. *Theoretical Computer Science*, 351(3):394–406, 2006. 1

[Mar11] Dániel Marx. Important separators and parameterized algorithms. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 5–10. Springer, 2011. 1

[Moi09] Ankur Moitra. Approximation algorithms for multicommodity-type problems with guarantees independent of the graph size. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 3–12. IEEE, 2009. 6

[MPRS20] Pranabendu Misra, Fahad Panolan, MS Ramanujan, and Saket Saurabh. Linear representation of transversal matroids and gammoids parameterized by rank. *Theoretical Computer Science*, 818:51–59, 2020. 6

[MR14] Daniel Marx and Igor Razgon. Fixed-parameter tractability of multicut parameterized by the size of the cutset. *SIAM Journal on Computing*, 43(2):355, 2014. 1, 2

[MSZ21] Jayakrishnan Madathil, Roohani Sharma, and Meirav Zehavi. A sub-exponential fpt algorithm and a polynomial kernel for minimum directed bisection on semicomplete digraphs. *Algorithmica*, 83:1861–1884, 2021. 4

[Räc08] Harald Räcke. Optimal hierarchical decompositions for congestion minimization in networks. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 255–264, 2008. 4

[SY22] Thatchaphol Saranurak and Sorrachai Yingchareonthawornchai. Deterministic small vertex connectivity in almost linear time. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 789–800. IEEE, 2022. 6

A Appendix

Lemma 1.2. *For any positive integer k, the following two statements hold.*

1. *There exist infinitely many positive integers c, such that for each c, there is a directed graph $G_{c,k}$ and disjoint subsets $S, T \subseteq V(G)$ of vertices with $|T| = c$ and $|S| = 1$ so that there are at least $\binom{|T|}{\leq k}$ A-B important separators of size at most k across all choices $A \subseteq S$, $B \subseteq T$.*
2. *There exist infinitely many positive integers c, such that for each c, there is a directed graph $G_{c,k}$ and disjoint subsets $S, T \subseteq V(G)$ of vertices with $|S| = c$ and $|T| = k+1$ so that there are at least $\binom{|S|}{\leq k}$ A-B important separators of size at most k across all choices $A \subseteq S$, $B \subseteq T$.*

Proof. Essentially, we will construct undirected graphs: these graphs can be made directed simply by adding for every undirected edge (u, v), the directed edges (u, v) and (v, u).

For the first part, consider the star on $c+1$ vertices for some $c \gg k$. Let S consist of the single center vertex, and let T consist of all the leaves. Observe that for any $T' \subseteq T$ of size at most k, the separator which consists of exactly the vertices T' is an important S-T' separator of size at most k. Thus there are at least $\binom{|T|}{\leq k}$ such separators.

For the second part, consider again a star, but this time we blow up the center vertex to a clique of size $k+1$. Formally, consider the graph G on $c+k+1$ vertices, where $T \subseteq V(G)$ is a clique of size $k+1$, and $S \subseteq V(G)$ forms the remaining c vertices. Each vertex of S is adjacent to each vertex of T, and not adjacent to any vertex of S. Thus the graph is a star with a core T of size $k+1$ and c leaves which form the set S. Now for every subset $S' \subseteq S$ with $|S'| \leq k$, S' itself forms an S'-T important separator of size at most k. Thus there are at least $\binom{|S|}{\leq k}$ such separators. □

Lemma A.1 (Tight example for important separator preservation)**.** *There is a graph G, source vertex $s \in V(G)$, integer k, sink vertices $B \subseteq V(G)$ with $|B| = 2k$ and an s-B important separator $X \subseteq V(G)$ with $|X| = k$, such that X is not an s-B' important separator for any $B' \subset B$.*

Proof. Consider the graph in Figure 2. X is an important s-B separator of size 3. However, clearly X is not an important s-B' important separator for any $B' \subset B$, since we must have $B' \cap \{b_{2i-1}, b_{2i}\} \leq 1$ for some $i \in [3]$ and this means that $(X \setminus \{u_i\}) \cup v_{2i-1}$ or $(X \setminus \{u_i\}) \cup v_{2i}$ is an s-B' separator closer to B' than X, contradicting the importance of X.

□

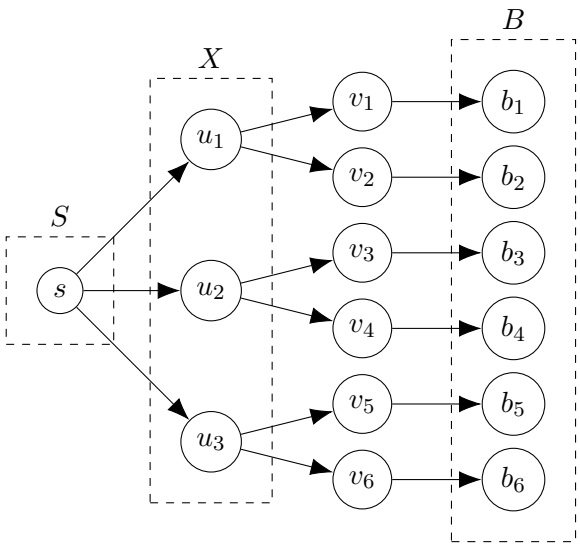

Figure 2: Tight example for important separator preservation.