SIAM J. COMPUT.
Vol. 42, No. 4, pp. 1674–1696

FIXED-PARAMETER TRACTABILITY OF DIRECTED MULTIWAY CUT PARAMETERIZED BY THE SIZE OF THE CUTSET*

RAJESH CHITNIS†, MOHAMMADTAGHI HAJIAGHAYI†, AND DÁNIEL MARX‡

Abstract. Given a directed graph G, a set of k terminals, and an integer p, the DIRECTED VERTEX MULTIWAY CUT problem asks whether there is a set S of at most p (nonterminal) vertices whose removal disconnects each terminal from all other terminals. DIRECTED EDGE MULTIWAY CUT is the analogous problem where S is a set of at most p edges. These two problems are indeed known to be equivalent. A natural generalization of the multiway cut is the MULTICUT problem, in which we want to disconnect only a set of k given pairs instead of all pairs. Marx [*Theoret. Comput. Sci.*, 351 (2006), pp. 394–406] showed that in undirected graphs VERTEX/EDGE MULTIWAY cut is fixed-parameter tractable (FPT) parameterized by p. Marx and Razgon [*Proceedings of the 43rd ACM Symposium on Theory of Computing*, 2011, pp. 469–478] showed that undirected MULTICUT is FPT and DIRECTED MULTICUT is W[1]-hard parameterized by p. We complete the picture here by our main result, which is that both DIRECTED VERTEX MULTIWAY CUT and DIRECTED EDGE MULTIWAY CUT can be solved in time $2^{2^{O(p)}} n^{O(1)}$, i.e., FPT parameterized by size p of the cutset of the solution. This answers an open question raised by the aforementioned papers. It follows from our result that DIRECTED EDGE/VERTEX MULTICUT is FPT for the case of $k = 2$ terminal pairs, which answers another open problem raised by Marx and Razgon.

Key words. multiway cut, fixed-parameter tractability, directed graphs

AMS subject classifications. 68W40, 05C85

DOI. 10.1137/12086217X

1. Introduction.

Ford and Fulkerson [11] gave the classical result on finding a minimum cut that separates two terminals s and t in 1956. A natural and well-studied generalization of the minimum $s-t$ cut problem is MULTIWAY CUT, in which, given a graph G and a set of terminals $\{s_1, s_2, \ldots, s_k\}$, the task is to find a minimum subset of vertices or edges whose deletion disconnects all the terminals from one another. Dahlhaus et al. [8] showed that the edge version in undirected graphs is APX-complete for $k \geq 3$. For the edge version Karger et al. [15] gave the current best known approximation ratio of 1.3438 for general k. The vertex version of the problem is known to be at least as hard as the edge version, and the current best approximation ratio is $2 - \frac{2}{k}$ [13].

The problem behaves very differently on directed graphs. Interestingly, for directed graphs, the edge and vertex versions turn out to be equivalent. Garg, Vazirani, and Yannakakis [13] showed that computing a minimum multiway cut in directed graphs is NP-hard and MAX SNP-hard already for $k = 2$. They also give an approx-

*Received by the editors January 11, 2012; accepted for publication (in revised form) April 22, 2013; published electronically August 6, 2013. A preliminary version of this paper appeared in *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, 2012 [5].
http://www.siam.org/journals/sicomp/42-4/86217.html

†Department of Computer Science, University of Maryland at College Park, College Park, MD 20742 (rchitnis@cs.umd.edu, hajiagha@cs.umd.edu). The research of the first and second authors was supported in part by NSF CAREER award 1053605, NSF grant CCF-1161626, ONR YIP award N000141110662, DARPA/AFOSR grant FA9550-12-1-0423, and a University of Maryland Research and Scholarship Award (RASA).

‡Computer and Automation Research Institute, Hungarian Academy of Sciences (MTA SZTAKI), Budapest, Hungary (dmarx@cs.bme.hu). This author's research was supported by the European Research Council (ERC) grant 280152.

TABLE 1
Summary of FPT results for UNDIRECTED MULTIWAY CUT. *Note that the O^* notation hides all factors which are polynomial in the size of the input.*

Problem	Running time	Paper
Vertex version	Nonconstructive FPT	Roberston and Seymour [24, 25]
	$O^*(4^{p^3})$	Marx [18]
	$O^*(4^p)$	Chen, Liu, and Lu [2]
	$O^*(4^p)$	Guillemot [14]
	$O^*(2^p)$	Cygan et al. [7]
Edge version	$O^*(2^p)$	Xiao [26]

imation algorithm with ratio 2 log k, which was later improved to ratio 2 by Naor and Zosin [20].

Rather than finding approximate solutions in polynomial time, one can look for exact solutions in time that is superpolynomial but still better than the running time obtained by brute force solutions. For example, Dahlhaus et al. [8] showed that undirected MULTIWAY CUT can be solved in time $n^{O(k)}$ on planar graphs, which can be an efficient solution if the number of terminals is small. On the other hand, on general graphs the problem becomes NP-hard already for $k = 3$. In both the directed and the undirected version, brute force can be used to check in time $n^{O(p)}$ whether a solution of size at most p exists: one can go through all sets of size at most p. Thus the problem can be solved in polynomial time if the optimum is assumed to be small. In the undirected case, significantly better running time can be obtained: the current fastest algorithms run in $O^*(2^p)$ time for both the vertex version [7] and the edge version [26] (the O^* notation hides all factors which are polynomial in size of input). That is, undirected MULTIWAY CUT is fixed-parameter tractable parameterized by the size of the cutset we remove. Recall that a problem is *fixed-parameter tractable* (FPT) with a particular parameter p if it can be solved in time $f(p)n^{O(1)}$, where f is an arbitrary function depending only on p; see [9, 10, 22] for more background. We give a brief summary of the race for faster FPT algorithms for UNDIRECTED MULTIWAY CUT in Table 1.

Our main result is that the directed version of MULTIWAY CUT is also FPT.

THEOREM 1.1 (main result). DIRECTED VERTEX MULTIWAY CUT *and* DIRECTED EDGE MULTIWAY CUT *can be solved in* $O^*(2^{2^{O(p)}})$ *time.*

Note that the hardness result of Garg, Vazirani, and Yannakakis [13] shows that in the directed case the problem is nontrivial (in fact, NP-hard) even for $k = 2$ terminals; our result holds without any bound on the number of terminals. The question was first asked explicitly in [18] and was also stated as an open problem in [19]. Our result shows in particular that directed multiway cut is solvable in polynomial time if the size of the optimum solution is $O(\log \log n)$, where n is the number of vertices in the digraph.

A more general problem is MULTICUT: the input contains a set $\{(s_1, t_1), \ldots, (s_k, t_k)\}$ of k pairs, and the task is to break every path from s_i to its corresponding t_i by the removal of at most p vertices. Very recently, it was shown that undirected MULTICUT is FPT parameterized by p [1, 19], but the directed version is unlikely to be FPT as it is W[1]-hard [19] with this parameterization. However, in the special case of $k = 2$ terminal pairs, there is a simple reduction from DIRECTED MULTICUT to DIRECTED MULTIWAY CUT; thus our result shows that the latter problem is FPT parameterized by p for $k = 2$. Let us briefly sketch the reduction. (Note that the reduction we sketch works only for the variant of DIRECTED MULTICUT which allows the deletion of terminals. Marx and Razgon [19] asked about the FPT status of this variant which is in fact equivalent to the one which does not allow deletion of the terminals.) Let (G, T, p) be a given instance of DIRECTED MULTICUT, and let $T = \{(s_1, t_1), (s_2, t_2)\}$. We construct an equivalent instance of DIRECTED MULTIWAY CUT as follows: graph G' is obtained by adding two new vertices s, t to the graph and adding the four edges $s \to s_1$, $t_1 \to t$, $t \to s_2$, and $t_2 \to s$. It is easy to see that the DIRECTED MULTIWAY CUT instance $(G', \{s, t\}, p)$ is equivalent to the original DIRECTED MULTICUT instance.[1]

COROLLARY 1.2. DIRECTED MULTICUT *with* $k = 2$ *can be solved in time* $O^*(2^{2^{O(p)}})$.

The complexity of the case $k = 3$ remains an interesting open problem.

Our techniques. Our algorithm for DIRECTED MULTIWAY CUT is inspired by the algorithm of Marx and Razgon [19] for undirected MULTICUT. In particular we use the technique of "random sampling of important separators" introduced in [19] and try to ensure that there is a solution whose "isolated part" is empty. However, DIRECTED MULTIWAY CUT behaves in a significantly different way than MULTICUT: at the same time, we are dealing with a much easier and a much harder situation. The first step in [19] is to reformulate the problem in such a way that the solution has to be a multiway cut of a certain set W of vertices; the technique of *iterative compression* allows us to reduce the original problem to this new version. As MULTIWAY CUT is already defined in terms of finding a multiway cut, this step is not necessary in our case. Furthermore, in [19], after ensuring that there is a solution whose "isolated part" is empty, the problem is reduced to ALMOST-2SAT. (Given a 2SAT formula and an integer k, is there an assignment satisfying all but k of the clauses?) This reduction works only if every component has at most two "legs"; a delicate branching algorithm is given to ensure this property. In the case of DIRECTED MULTIWAY CUT, the situation is much simpler: if there is a solution whose "isolated part" is empty, then the problem can be reduced to the undirected version, and then we can use the current fastest undirected algorithm [7], which runs in $O^*(2^p)$ time.

On the other hand, the fact that we are dealing with a directed graph makes the problem significantly harder (recall that DIRECTED MULTICUT is W[1]-hard parameterized by p; thus it is expected that not every undirected argument generalizes to the directed case). After defining a proper notion of directed important separators, the nontrivial interaction among two kinds of "shadows" forces us to do the random sampling of important separators in two independent steps, and the analysis becomes more delicate.

Independent and follow-up work. The fixed-parameter tractability of MULTICUT in undirected graphs parameterized only by the size of the cutset was shown independently by Marx and Razgon [19] and Bousquet, Daliault, and Thomassé [1]. Marx and Razgon [19] also showed that DIRECTED MULTICUT is W[1]-hard parameterized by the size of the cutset. The technique of random sampling of important separators introduced in [19] is a crucial element of our algorithm. A very different application of this technique was given by Lokshtanov and Marx [17] in the context of clustering problems.

[1] G has an $s_i \to t_i$ path for some i if and only if G' has an $s \to t$ or $t \to s$ path. This is because G has an $s_1 \to t_1$ path if and only if G' has an $s \to t$ path, and G has an $s_2 \to t_2$ path if and only if G' has a $t \to s$ path. This property of paths also holds after removing some vertices/edges, and thus the two instances are equivalent.

The preliminary version of this paper adapted the framework of random sampling of important separators to directed graphs and showed the fixed-parameter tractability of Directed Multiway Cut parameterized by the size of the cutset. This framework was later used by Kratsch et al. [16] to show the fixed-parameter tractability of Directed Multicut on directed acyclic graphs and by Chitnis et al. [4] to show the fixed-parameter tractability of Subset Directed Feedback Vertex Set. The latter paper improved the randomized sampling process to make the algorithms more efficient; in particular, this improvement results in an $O^*(2^{O(p^2)})$ algorithm for Directed Multiway Cut. The question of existence of a polynomial kernel for Directed Multiway Cut was answered negatively by Cygan et al. [6], who showed that Directed Multiway Cut (even for two terminals) does not have a polynomial kernel unless NP $\subseteq$ coNP/poly and the polynomial hierarchy collapses to the third level. An interesting open question is the complexity of Directed Multicut for $k = 3$ or with combined parameters k and p.

2. Preliminaries. A multiway cut is a set of edges/vertices that separate the terminal vertices from each other.

Definition 2.1 (multiway cut). *Let G be a directed graph, and let $T = \{t_1, t_2, \ldots, t_k\} \subseteq V(G)$ be a set of terminals.*

1. *$S \subseteq V(G)$ is a* vertex multiway cut *of (G, T) if $G \setminus S$ does not have a path from t_i to t_j for any $i \neq j$.*
2. *$S \subseteq E(G)$ is a* edge multiway cut *of (G, T) if $G \setminus S$ does not have a path from t_i to t_j for any $i \neq j$.*

In the edge case, it is straightforward to define the problem that we want to solve, as follows.

Directed Edge Multiway Cut
Input: A directed graph G, an integer p, and a set of terminals T.
Output: A multiway cut $S \subseteq E(G)$ of (G, T) of size at most p or "NO" if such a multiway cut does not exist.

In the vertex case, there is a slight technical issue in the definition of the problem: are the terminal vertices allowed to be deleted? We focus here on the version of the problem where the vertex multiway cut we are looking for has to be disjoint from the set of terminals. More generally, we define the problem in such a way that the graph has some *distinguished* vertices which cannot be included as part of any separator (and we assume that every terminal is a distinguished vertex). This can be modeled by considering weights on the vertices of the graph: weight of ∞ on each distinguished vertex and 1 on every nondistinguished vertex. We look only for solutions of finite weight. From here on, for a graph G we will denote by $V^\infty(G)$ the set of distinguished vertices of G with the meaning that these distinguished vertices cannot be part of any separator; i.e., all separators we consider are of finite weight. In fact, for any separator we can talk interchangeably about size or weight as these notions are the same since each vertex of separator has weight 1.

The main focus of the paper is the following vertex version, where we require $T \subseteq V^\infty(G)$; i.e., terminals cannot be deleted.

Directed Vertex Multiway Cut
Input: A directed graph G, an integer p, a set of terminals T and a set $V^\infty \supseteq T$ of distinguished vertices.
Output: A multiway cut $S \subseteq V(G) \setminus V^\infty(G)$ of (G, T) of size at most p or "NO" if such a multiway cut does not exist.

We note that if we want to allow the deletion of the terminal vertices, then it is not difficult to reduce the problem to the version defined above. For each terminal t we introduce a new vertex t', and we add the directed edges (t,t') and (t',t). Let the new graph be G', and let $T' = \{t' \mid t \in T\}$. Then there is a clear bijection between vertex multiway cuts which can include terminals in the instance (G,T,p) and vertex multiway cuts which cannot include terminals in the instance (G',T',p).

The two versions Directed Vertex Multiway Cut and Directed Edge Multiway Cut defined above are known to be equivalent. For the sake of completeness, we prove the equivalence in section 2.1. In the remaining part of the paper, we concentrate on finding an FPT algorithm for Directed Vertex Multiway Cut, which we henceforth call Directed Multiway Cut for brevity.

2.1. Equivalence of vertex and edge versions of Directed Multiway Cut. We first show how to solve the vertex version using the edge version. Let (G,T,p) be a given instance of Directed Vertex Multiway Cut, and let $V^{\infty}(G)$ be the set of distinguished vertices. We construct an equivalent instance (G',T',p) of Directed Edge Multiway Cut as follows. Let the set V' contain two vertices v^{in}, v^{out} for every $v \in V(G) \setminus V^{\infty}(G)$ and a single vertex $u^{\text{in}} = u^{\text{out}}$ for every $u \in V^{\infty}(G)$. The idea is that all incoming/outgoing edges of v in G will now be incoming/outgoing edges of v^{in} and v^{out}, respectively. For every vertex $v \in V(G) \setminus V^{\infty}(G)$, add an edge $(v^{\text{in}}, v^{\text{out}})$ to G'. Let us call these Type I edges. For every edge $(x,y) \in E(G)$, add $(p+1)$ parallel $(x^{\text{out}}, y^{\text{in}})$ edges. Let us call these Type II edges. Define $T' = \{v^{\text{in}} \mid v \in T\}$. Note that the number of terminals is preserved. We have the following lemma.

Lemma 2.2. *(G,T,p) is a yes-instance of* Directed Vertex Multiway Cut *if and only if (G',T',p) is a yes-instance of* Directed Edge Multiway Cut.

Proof. Suppose G has a vertex multiway cut, say S, of size at most p. Then the set $S' = \{(v^{\text{in}}, v^{\text{out}}) \mid v \in S\}$ is clearly an edge multiway cut for G' and $|S'| = |S| \leq p$.

Suppose G' has an edge multiway cut, say S', of size at most p. Note that it does not help to pick in S any edges of Type II as each edge has $(p+1)$ parallel copies and our budget is p. So let $S = \{v \mid (v^{\text{in}}, v^{\text{out}}) \in S'\}$. Then S is a vertex multiway cut for G and $|S| \leq |S'| \leq p$. ∎

We now show how to solve the edge version using the vertex version. Let (G,T,p) be a given instance of Directed Edge Multiway Cut. We construct an equivalent instance (G',T',p) of Directed Vertex Multiway Cut as follows. For each vertex $u \in V(G) \setminus T$, create a set C_u which contains u along with p other copies of u. For $t \in T$ we let $C_t = \{t\}$. For each edge $(u,v) \in E(G)$ create a vertex β_{uv}. Add edges (x, β_{uv}) for all $x \in C_u$ and (β_{uv}, y) for all $y \in C_v$. Define $T' = \bigcup_{t \in T} C_t = T$. Let $V^{\infty}(G') = T'$

Lemma 2.3. *(G,T,p) is a yes-instance of* Directed Edge Multiway Cut *if and only if (G',T',p) is a yes-instance of* Directed Vertex Multiway Cut.

Proof. Suppose G has an edge multiway cut, say S, of size at most p. Then the set $S' = \{\beta_{uv} \mid (u,v) \in S\}$ is clearly a vertex multiway cut for G' and $|S'| = |S| \leq p$.

Suppose G' has a vertex multiway cut, say S', of size at most p. Note that it does not help to pick in S any vertices from the C_z of any vertex $z \in V(G) \setminus T$ as each vertex has $(p+1)$ equivalent copies and our budget is p. So let $S = \{(u,v) \mid \beta_{uv} \in S'\}$. Then S is a edge multiway cut for G and $|S| \leq |S'| \leq p$. ∎

2.2. Separators and shadows. The crucial idea in the algorithm of [19] for (the vertex version of) undirected Multicut is to get rid of the "isolated part" of

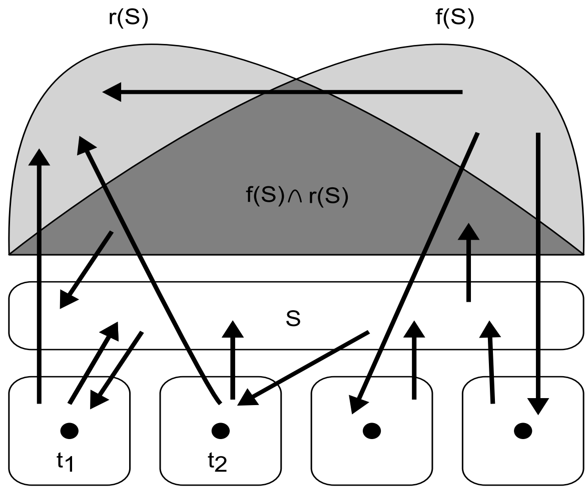

FIG. 1. *For every vertex $v \in f(S)$, the set S is a $T-v$ separator. For every vertex $w \in r(S)$, the set S is a $w-T$ separator. For every vertex $y \in f(S) \cap r(S)$, the set S is both a $T-y$ and a $y-T$ separator. Finally for every $z \in V(G) \setminus [S \cup r(S) \cup f(S) \cup T]$, there are both $z-T$ and $T-z$ paths in the graph $G \setminus S$. Note that every such vertex z belongs to a strongly connected component of $G \setminus S$ containing T and there are no edges between these components.*

the solution S. We use a similar concept here, but we use the term *shadow*, as it is more expressive for directed graphs.

DEFINITION 2.4 (separator). *Let G be a directed graph, and let $V^\infty(G) \supseteq T$ be the set of distinguished ("undeletable") vertices. Given two disjoint nonempty sets $X, Y \subseteq V$, we call a set $S \subseteq V \setminus (X \cup Y \cup V^\infty)$ an $X-Y$* separator *if there is no path from X to Y in $G \setminus S$. A set S is a* minimal $X-Y$ separator *if no proper subset of S is an $X-Y$ separator.*

Note that here we explicitly define the $X-Y$ separator S to be disjoint from X and Y.

DEFINITION 2.5 (shadows). *Let G be a graph, and let T be a set of terminals. Let $S \subseteq V(G) \setminus V^\infty(G)$ be a subset of vertices.*

1. *The* forward shadow *$f_{G,T}(S)$ of S (with respect to T) is the set of vertices v such that S is a $T-\{v\}$ separator in G.*
2. *The* reverse shadow *$r_{G,T}(S)$ of S (with respect to T) is the set of vertices v such that S is a $\{v\}-T$ separator in G.*

The shadow *of S (with respect to T) is the union of $f_{G,T}(S)$ and $r_{G,T}(S)$.*

That is, we can imagine T as a light source with light spreading on the directed edges. The forward shadow is the set of vertices that remain dark if the set S blocks the light, hiding v from T's sight. In the reverse shadow, we imagine that light is spreading backwards on the edges. We abuse the notation slightly and write $v-T$ separator instead of $\{v\}-T$ separator. We also drop G and T from the subscript if they are clear from the context. Note that S itself is not in the shadow of S (as, by definition, a $T-v$ or $v-T$ separator needs to be disjoint from T and v); that is, S and $f_{G,T}(S) \cup r_{G,T}(S)$ are disjoint. See Figure 1 for an illustration.

3. Overview of our algorithm. We say that a solution S of DIRECTED MULTIWAY CUT is *shadowless* (with respect to T) if $f(S) = r(S) = \emptyset$. The following lemma shows the importance of *shadowless solutions* for DIRECTED MULTIWAY CUT.

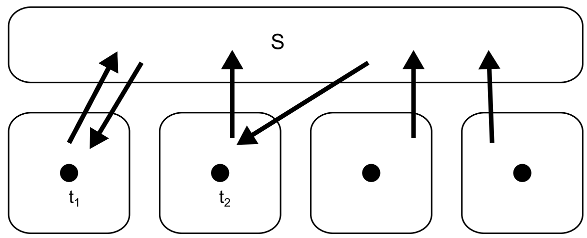

FIG. 2. *A shadowless solution S for a* DIRECTED MULTIWAY CUT *instance. Every vertex of $G\setminus S$ is in the strongly connected component of some terminal t_i. There are no edges between the strongly connected components of the terminals t_i; thus S is also a solution of the underlying* UNDIRECTED MULTIWAY CUT *instance.*

Clearly, any solution of the underlying undirected instance (where we disregard the orientation of the edges) is a solution for DIRECTED MULTIWAY CUT cut. The converse is not true in general: a solution of the directed problem is not always a solution of the undirected problem. However, the following lemma shows that the converse statement is true for shadowless solutions of the directed instance.

LEMMA 3.1. *Let G^* be the underlying undirected graph of G. If S is a* shadowless *solution for an instance (G, T, p) of* DIRECTED MULTIWAY CUT*, then S is also a solution for the instance (G^*, T, p) of* UNDIRECTED MULTIWAY CUT*.*

Proof. If S is a shadowless solution, then for each vertex v in $G \setminus S$, there is a $t_1 \to v$ path and a $v \to t_2$ path for some $t_1, t_2 \in T$. As S is a solution, it is not possible that $t_1 \neq t_2$: this would give a $t_1 \to t_2$ path in $G \setminus S$. Therefore, if S is a shadowless solution, then each vertex in the graph $G \setminus S$ belongs to the strongly connected component of exactly one terminal. A directed edge between the strongly connected components of t_i and t_j would imply the existence of either a $t_i \to t_j$ or a $t_j \to t_i$ path, which contradicts the fact that S is a solution of the DIRECTED MULTIWAY CUT instance. Hence the strongly connected components of $G \setminus S$ are exactly the same as the weakly connected components of $G \setminus S$; i.e., S is also a solution for the underlying instance of UNDIRECTED MULTIWAY CUT. ∎

An illustration of Lemma 3.1 is given in Figure 2. Lemma 3.1 shows that if we can transform the instance in a way that ensures the existence of a shadowless solution, then we can reduce the problem to undirected MULTIWAY CUT and use the $O^*(4^p)$ algorithm for that problem due to Guillemot [14] which can handle the case when there are some distinguished vertices similar to what we consider. Our transformation is based on two ingredients: random sampling of important separators and reduction of the instance using the `torso` operation. These techniques were introduced by Marx and Razgon [19] for the undirected MULTICUT problem. In section 4, we review these tools and adapt them for directed graphs.

Random sampling of important separators. As a first step in reducing the problem to a shadowless instance, we need a set Z that has the following property:

(*) There is a solution S^* such that Z contains the shadow of S^*, but Z is disjoint from S^*.

If we have a set Z that satisfies property (*), we modify the instance in a way that removes the set Z. The modification is done such that S^* remains a solution of the reduced instance; in fact, it becomes a shadowless solution. This means that the problem can be solved by Lemma 3.1. This process of getting rid of the set Z in an appropriate way is accomplished by the `torso` operation defined below.

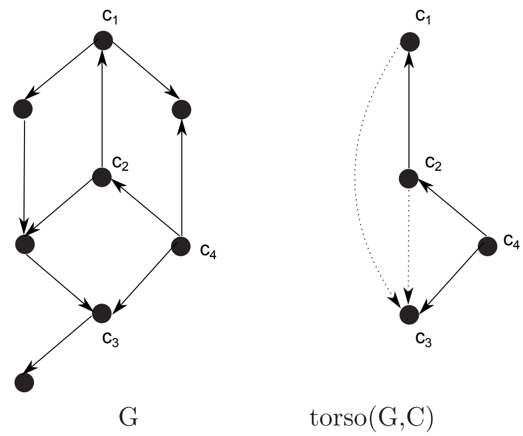

FIG. 3. *Let $C = \{c_1, c_2, c_3, c_4\}$. In the graph* $\texttt{torso}(G, C)$ *the edges (c_4, c_3) and (c_4, c_2) carry over from G. The new edges (shown by dotted arrows) that get added because of the* `torso` *operation are (c_1, c_3) and (c_2, c_3).*

Unfortunately, when we are trying to construct the set Z, we do not know anything about the solutions of the instance, and in particular we have no way of checking whether a given set Z satisfies property (*). Nevertheless, we use a randomized procedure that creates a set Z, and we give a lower bound on the probability that Z satisfies property (*). For the construction of this set Z, we use a very specific probability distribution that was introduced in [19]. This probability distribution is based on randomly selecting "important separators" and taking the union of their shadows. At this point, we can consider the sampling as a black-box function "RandomSet(G, T, p)" that returns a random subset $Z \subseteq V(G)$ according to a probability distribution that satisfies certain properties. The precise description of this function and the properties of the distribution it creates is described in section 4.2 (see Theorem 4.10). The randomized selection can be derandomized: the randomized selection can be turned into a deterministic algorithm that returns a bounded number of sets such that at least one of them satisfies the required property (section 4.3). To make the description of the algorithm simpler, we focus on the randomized algorithm in this section.

Torsos. We use the function RandomSet(G, T, p) to construct a set Z of vertices that we want to get rid of. However we must be careful: when getting rid of the set Z we should ensure that the information relevant to Z is captured in the reduced instance. This is exactly accomplished by the `torso` operation which removes a set of vertices without making the problem any easier. We formally define this operation as follows.

Definition 3.2 (`torso`). *Let G be a directed graph, and let $C \subseteq V(G)$. The graph* $\texttt{torso}(G, C)$ *has vertex set C, and there is a (directed) edge (a, b) in* $\texttt{torso}(G, C)$ *if there is an $a \to b$ path in G whose internal vertices are not in C.*

See Figure 3 for an example of the `torso` operation. Note that if $a, b \in C$ and (a, b) is a directed edge of G, then $\texttt{torso}(G, C)$ contains (a, b) as well. Thus $G[C]$, which is the graph induced by C in G, is a subgraph of $\texttt{torso}(G, C)$. The following lemma shows that the `torso` operation preserves separation inside C.

Lemma 3.3 (`torso` preserves separation). *Let G be a directed graph, and let $C \subseteq V(G)$. Let $G' =$ $\texttt{torso}(G, C)$ and $S \subseteq C$. For $a, b \in C \setminus S$, the graph $G \setminus S$ has an $a \to b$ path if and only if $G' \setminus S$ has an $a \to b$ path.*

Proof. Let P be a path from a to b in G. Suppose P is disjoint from S. Then P contains vertices from C and $V(G) \setminus C$. Let u, v be two vertices of C such that every vertex of P between u and v is from $V(G) \setminus C$. Then by definition there is an edge (u, v) in $\texttt{torso}(G, C)$. Using such edges, we can modify P to obtain an $a \to b$ path that lies completely in $\texttt{torso}(G, C)$ but avoids S.

Conversely, suppose that P' is an $a \to b$ path in $\texttt{torso}(G, C)$ and avoids $S \subseteq C$. If P' uses an edge $(u, v) \notin E(G)$, then this means that there is a $u \to v$ path P'' whose internal vertices are not in C. Using such paths, we modify P to get an $a \to b$ path P_0 that uses only edges from G. Since $S \subseteq C$, we have that the new vertices on the path are not in S and so P_0 avoids S. ∎

If we want to remove a set Z of vertices, then we create a new instance by taking the torso on the *complement* of Z.

Definition 3.4. *Let $I = (G, T, p)$ be an instance of* Directed Multiway Cut *and $Z \subseteq V(G) \setminus T$. The reduced instance $I/Z = (G', T', p)$ is defined as*

- $G' = \texttt{torso}(G, V(G) \setminus Z)$,
- $T' = T$.

The following lemma states that the operation of taking the torso does not make the Directed Multiway Cut problem easier for any $Z \subseteq V(G) \setminus T$ in the sense that any solution of the reduced instance I/Z is a solution of the original instance I. Moreover, if we perform the torso operation for a Z that is large enough to contain the shadow of some solution S^* but at the same time small enough to be disjoint from S^*, then S^* remains a solution for the reduced instance I/Z and in fact is a shadowless solution for I/Z. Therefore, our goal is to randomly select a set Z in such a way that we can bound the probability that Z satisfies property (*) defined above for some hypothetical solution S^*.

Lemma 3.5 (creating a shadowless instance). *Let $I = (G, T, p)$ be an instance of* Directed Multiway Cut *and $Z \subseteq V(G) \setminus T$.*

1. *If S is a solution for I/Z, then S is also a solution for I.*
2. *If S is a solution for I such that $f_{G,T}(S) \cup r_{G,T}(S) \subseteq Z$ and $S \cap Z = \emptyset$, then S is a shadowless solution for I/Z.*

Proof. Let G' be the graph $\texttt{torso}(G, V(G) \setminus Z)$. To prove the first part, suppose that $S \subseteq V(G')$ is a solution for I/Z and S is not a solution for I. Then there are terminals $t_1, t_2 \in T$ such that there is a $t_1 \to t_2$ path P in $G \setminus S$. As $t_1, t_2 \in T$ and $Z \subseteq V(G) \setminus T$, we have that $t_1, t_2 \in V(G) \setminus Z$. In fact, we have $t_1, t_2 \in (V(G) \setminus Z) \setminus S$. Lemma 3.3 implies that there is a $t_1 \to t_2$ path in $G' \setminus S$, which is a contradiction as S is a solution for I/Z.

For the second part of the lemma, let S be a solution for I such that $S \cap Z = \emptyset$ and $f_{G,T}(S) \cup r_{G,T}(S) \subseteq Z$. We want to show that S is a shadowless solution for I/Z. First we show that S is a solution for I/Z. Suppose to the contrary that there are terminals $x', y' \in T'(= T)$ such that $G' \setminus S$ has an $x' \to y'$ path. As $x', y' \in V(G) \setminus Z$, Lemma 3.3 implies that $G \setminus S$ also has an $x' \to y'$ path, which is a contradiction as S is a solution of I.

Finally, we show that S is shadowless in I/Z; i.e., $r_{G',T}(S) = \emptyset = f_{G',T}(S)$. We prove only that $r_{G',T}(S) = \emptyset$: the argument for $f_{G',T}(S) = \emptyset$ is analogous. Assume to the contrary that there exists $w \in r_{G',T}(S)$ (note that we have $w \in V(G')$, i.e., $w \notin Z$). So S is a $w - T$ separator in G'; i.e., there is no $w - T$ path in $G' \setminus S$. Lemma 3.3 gives that there is no $w - T$ path in $G \setminus S$; i.e., $w \in r_{G,T}(S)$. But $r_{G,T}(S) \subseteq Z$, and so we have $w \in Z$, which is a contradiction. Thus $r_{G,T}(S) \subseteq Z$ in G implies that $r_{G',T}(S) = \emptyset$. ∎

Algorithm 1. FPT Algorithm for Directed Multiway Cut.

Input: An instance $I_1 = (G_1, T, p)$ of Directed Multiway Cut.

1: Let $Z_1 = \textsc{RandomSet}(G_1, T, p)$.
2: Let $G_2 = (G_1)_{\text{rev.}}$ {Reverse the orientation of every edge.}
3: Let $V^\infty(G_2) = V^\infty(G_1) \cup Z_1$. {Set weight of every vertex of Z_1 to ∞.}
4: Let $Z_2 = \textsc{RandomSet}(G_2, T, p)$.
5: Let $Z = Z_1 \cup Z_2$.
6: Let $G_3 = \texttt{torso}(G_1, V(G) \setminus Z)$. {Get rid of Z.}
7: Solve the underlying *undirected* instance (G_3^*, T, p) of Multiway Cut.
8: **if** (G_3^*, T, p) has a solution S **then**
9: **return** S
10: **else**
11: **return** "NO"

The algorithm. The description of our algorithm is given in Algorithm 1. Recall that we are trying to solve a version of Directed Multiway Cut where we are given a set V^∞ of distinguished vertices which are undeletable, i.e., have infinite weight.

Due to the delicate way separators behave in directed graphs, we construct the set Z in two phases, calling the function RandomSet twice. Our aim is to show that there is a solution S such that we can give a lower bound on the probability that Z_1 contains $r_{G_1,T}(S)$ and Z_2 contains $f_{G_1,T}(S)$. Note that the graph G_2 obtained in step 2 depends on the set Z_1 returned in step 1 (as we made the weight of every vertex in Z_1 infinite); thus the distribution of the second random sampling depends on the result Z_1 of the first random sampling. This means that we cannot make the two calls in parallel.

We use the `torso` operation to remove the vertices in $Z = Z_1 \cup Z_2$ (step 5), and then solve the undirected Multiway Cut instance obtained by disregarding the orientation of the edges. For this purpose, we can use the algorithm of Guillemot [14] that solves the undirected problem in time $O^*(4^p)$. Note that the algorithm for undirected Multiway Cut in [14] explicitly considers the variant where we have a set of distinguished vertices which cannot be deleted.

The following two lemmas show that Algorithm 1 is a correct randomized algorithm. One direction is easy to see: the algorithm has no false positives.

Lemma 3.6. *Let $I_1 = (G_1, T, p)$ be an instance of* Directed Multiway Cut*. If Algorithm* 1 *returns a set S, then S is a solution for I_1.*

Proof. Any solution S of the undirected instance (G_3^*, T, p) returned by Algorithm 1 is clearly a solution of the directed instance (G_3, T, p) as well. By Lemma 3.5(1) the `torso` operation does not make the problem easier by creating new solutions. Hence S is also a solution for $I_1 = (G_1, T, p)$. ∎

The following lemma shows that if the instance has a solution, then the algorithm finds one with certain probability.

Lemma 3.7. *Let $I_1 = (G_1, T, p)$ be an instance of* Directed Multiway Cut*. If I_1 is a yes-instance of* Directed Multiway Cut*, then Algorithm* 1 *returns a set S which is a solution for I with probability at least $2^{-2^{O(p)}}$.*

By Lemma 3.5(2), we can prove Lemma 3.7 by showing that if I_1 is a yes-instance, then there exists a solution S^* such that Z satisfies the two requirements $Z \cap S = \emptyset$ and $f_{G_1,T}(S) \cup r_{G_1,T}(S) \subseteq Z$ with suitable probability. This requires a deeper analysis of the structure of optimum solutions and the probability distribution behind the function RandomSet(G, T, p). Hence we defer the proof of Lemma 3.7 to section 5.

Derandomization. In section 4.3, we present a deterministic variant of RandomSet(G, T, p), which, instead of returning a random set Z, returns a deterministic set $Z_1, \ldots, Z_t$ of $O^*(2^{2^{O(p)}})$ sets. Instead of bounding the probability that the random set Z has the required property with some probability, we prove that at least one Z_i always satisfies the property. Therefore, in steps 1 and 3 of Algorithm 1, we can replace RandomSet with this deterministic variant and branch on the choice of one Z_i from the returned sets. By the properties of the deterministic algorithm, if I_1 is a yes-instance, then Z has property (*) in at least one of the branches and therefore the algorithm finds a correct solution for I_1. The branching increases the running time only by a factor of $(O^*(2^{2^{O(p)}}))^2$ and therefore the total running time is $O^*(2^{2^{O(p)}})$.

4. Important separators and random sampling. This section reviews the notion of important separators and the random sampling technique introduced by Marx and Razgon [19]. As [19] used these concepts for undirected graphs and we need them for directed graphs, we give a self-contained presentation without relying on earlier work.

4.1. Important separators. Marx [18] introduced the concept of *important separators* to deal with the Undirected Multiway Cut problem. Since then it has been used implicitly or explicitly in, e.g., [2, 3, 17, 19, 23] in the design of fixed-parameter algorithms. In this section, we define and use this concept in the setting of directed graphs. Roughly speaking, an important separator is a separator of small size that is *maximal* with respect to the set of vertices on one side.

Definition 4.1 (important separator). *Let G be a directed graph, and let $X, Y \subseteq V$ be two disjoint nonempty sets. A minimal $X - Y$ separator S is called an* important $X - Y$ separator *if there is no $X - Y$ separator S' with $|S'| \leq |S|$ and $R^+_{G \setminus S}(X) \subset R^+_{G \setminus S'}(X)$, where $R^+_A(X)$ is the set of vertices reachable from X in A.*

Let X, Y be disjoint sets of vertices of an *undirected graph*. Then for every $p \geq 0$ it is known [2, 18] that there are at most 4^p important $X - Y$ separators of size at most p for any sets X, Y. The next lemma shows that the same bound holds for important separators even in directed graphs.

Lemma 4.2 (number of important separators). *Let $X, Y \subseteq V(G)$ be disjoint sets in a* directed *graph G. Then for every $p \geq 0$ there are at most 4^p important $X - Y$ separators of size at most p. Furthermore, we can enumerate all these separators in time $O(4^p \cdot p(|V(G) + |E(G)|))$.*

The proof of Lemma 4.2 is long and follows the same techniques as the proof in undirected graphs (see, e.g., [19, 17]). Therefore, it is deferred to Appendix A to maintain the flow of the main result. For ease of notation, we now define the following collection of important separators.

Definition 4.3. *Given an instance (G, T, p) of* Directed Multiway Cut*, the set $\mathcal{I}_p$ contains the set $S \subseteq V(G)$ if S is an important $v - T$ separator of size at most p in G for some vertex v in $V(G) \setminus T$.*

Remark 4.4. It follows from Lemma 4.2 that $|\mathcal{I}_p| \leq 4^p \cdot |V(G)|$ and we can enumerate the sets in $\mathcal{I}_p$ in time $O^*(4^p)$.

We now define a special type of shadows which we use later for the random sampling.

Definition 4.5 (exact shadows). *Let G be a directed graph and $T \subseteq V(G)$ a set of terminals. Let $S \subseteq V(G) \setminus V^\infty(G)$ be a set of vertices. Then for $v \in V(G)$ we say*

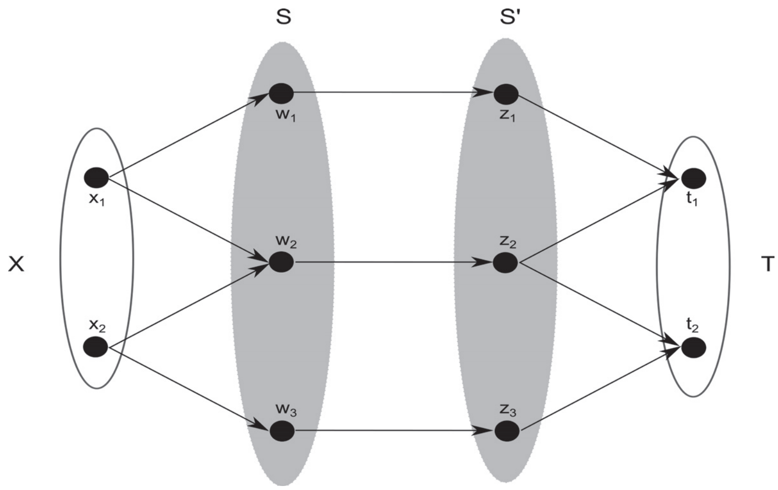

FIG. 4. *S is a minimal $X-Y$ separator, but it is not an important $X-T$ separator as S' satisfies $|S'|=|S|$ and $R^+_{G\setminus S}(X)=X\subset X\cup S=R^+_{G\setminus S'}(X)$. In fact it is easy to check that the only important $X-T$ separator of size 3 is S'. If $p\geq 2$, then the set $\{z_1,z_2\}$ is in $\mathcal{I}_p$ since it is an important x_1-T separator of size 2. Finally, x_1 belongs to the "exact reverse shadow" of each of the sets $\{w_1,w_2\},\{w_1,z_2\},\{w_2,z_1\}$, and $\{z_1,z_2\}$ since they are all minimal x_1-T separators. However x_1 does not belong to the exact reverse shadow of the set S as it is not a minimal x_1-T separator.*

that

1. *v is in the "exact reverse shadow" of S (with respect to T) if S is a minimal $v-T$ separator in G, and*
2. *v is in the "exact forward shadow" of S (with respect to T) if S is a minimal $T-v$ separator in G.*

We refer the reader to Figure 4 for examples of Definitions 4.1, 4.3, and 4.5. The exact reverse shadow of S is a subset of the reverse shadow of S: it contains a vertex v only if every vertex $w\in S$ is "useful" in separating v, that is, vertex w can be reached from v, and T can be reached from w. This slight difference between the shadow and the exact shadow will be crucial in the analysis of the algorithm (see section 5 and Remark 4.8).

The random sampling described in section 4.2 (Theorem 4.10) randomly selects members of $\mathcal{I}_p$ and creates a subset of vertices by taking the union of the exact reverse shadows of the selected separators. The following lemma will be used to give an upper bound on the probability that a vertex is covered by the union.

LEMMA 4.6. *Let z be any vertex. Then there are at most 4^p members of $\mathcal{I}_p$ which contain z in their exact reverse shadows.*

For the proof of Lemma 4.6, first we need to establish the following.

LEMMA 4.7. *If $S\in\mathcal{I}_p$ and v is in the exact reverse shadow of S, then S is an important $v-T$ separator.*

Proof. Let w be the witness that S is in $\mathcal{I}_p$, i.e., S is an important $w-T$ separator in G. Let v be any vertex in the exact reverse shadow of S, which means that S is a minimal $v-T$ separator in G. Suppose that S is not an important $v-T$ separator. Then there exists a $v-T$ separator S' such that $|S'|\leq|S|$ and $R^+_{G\setminus S}(v)\subset R^+_{G\setminus S'}(v)$. We will arrive to a contradiction by showing that $R^+_{G\setminus S}(w)\subset R^+_{G\setminus S'}(w)$, i.e., S is not an important $w-T$ separator.

First, we claim that S' is an $(S \setminus S') - T$ separator. Suppose that there is a path P from some $x \in S \setminus S'$ to T that is disjoint from S'. As S is a minimal $v - T$ separator, there is a path Q from v to x whose internal vertices are disjoint from S. Furthermore, $R^+_{G \setminus S}(v) \subset R^+_{G \setminus S'}(v)$ implies that the internal vertices of Q are disjoint from S' as well. Therefore, concatenating Q and P gives a path from v to T that is disjoint from S', contradicting the fact that S' is a $v - T$ separator.

We show that S' is a $w-T$ separator and its existence contradicts the assumption that S is an important $w - T$ separator. First we show that S' is a $w - T$ separator. Suppose that there is a $w - T$ path P disjoint from S'. Path P has to go through a vertex $y \in S \setminus S'$ (as S is a $w-T$ separator). Thus by the previous claim, the subpath of P from y to T has to contain a vertex of S', a contradiction.

Finally, we show that $R^+_{G \setminus S}(w) \subseteq R^+_{G \setminus S'}(w)$. As $S \neq S'$ and $|S'| \leq |S|$, this will contradict the assumption that S is an important $w - T$ separator. Suppose that there is a vertex $z \in R^+_{G \setminus S}(w) \setminus R^+_{G \setminus S'}(w)$, and consider a $w - z$ path that is fully contained in $R^+_{G \setminus S}(v)$, i.e., disjoint from S. As $z \notin R^+_{G \setminus S'}(v)$, path Q contains a vertex $q \in S' \setminus S$. Since S' is a minimal $v - T$ separator, there is a $v - T$ path that intersects S' only in q. Let P be the subpath of this path from q to T. If P contains a vertex $r \in S$, then the subpath of P from r to T contains no vertex of S' (as $z \neq r$ is the only vertex of S' on P), contradicting our earlier claim that S' is an $(S \setminus S') - T$ separator. Thus P is disjoint from S, and hence the concatenation of the subpath of Q from w to q and the path P is a $w-T$ path disjoint from S, a contradiction. ∎

Lemma 4.6 easily follows from Lemma 4.7. Let J be a member of $\mathcal{I}_p$ such that z is in the exact reverse shadow of J. By Lemma 4.7, J is an important $z-T$ separator. By Lemma 4.2, there are at most 4^p important $z - T$ separators of size at most p, and so z belongs to at most 4^p exact reverse shadows.

Remark 4.8. It is crucial to distinguish between "reverse shadow" and "exact reverse shadow": Lemma 4.7 (and hence Lemma 4.6) does not remain true if we remove the word "exact." Consider the following example (see Figure 5). Let a_1, $\ldots$, a_r be vertices such that there is an edge going from every a_i to every vertex of $T = \{t_1, t_2, \ldots, t_k\}$. For every $1 \leq i \leq r$, let b_i be a vertex with an edge going from b_i to a_i. For every $1 \leq i < j \leq r$, let $c_{i,j}$ be a vertex with two edges going from $c_{i,j}$ to a_i and a_j. Then every set $\{a_i, a_j\}$ is in $\mathcal{I}_p$, since it is an important $c_{i,j} - T$ separator. This means that every b_i is in the reverse shadow of $r - 1$ members of $\mathcal{I}_p$, namely the sets $\{a_j, ai_i\}$ for $1 \leq i \neq j \leq r$. However, b_i is in the *exact* reverse shadow of exactly one member of $\mathcal{I}_p$, the set $\{a_i\}$.

4.2. Random sampling. In this section, we adapt the random sampling of [19] to directed graphs. We try to present it in a self-contained way that might be useful for future applications.

Roughly speaking, we want to select a random set Z such that for every pair (S, Y) where Y is in the reverse shadow of S, the probability that Z is disjoint from S but contains Y can be bounded from below. We can guarantee such a lower bound only if (S, Y) satisfies two conditions. First, it is not enough that Y is in the shadow of S (or in other words, S is an $Y - T$ separator), but S should contain important separators separating the vertices of Y from T (see Theorem 4.10 for the exact statement). Second, a vertex of S cannot be in the reverse shadow of other vertices of S; this is expressed by the following technical definition.

Definition 4.9 (thin). *Let G be a directed graph and $T \subseteq V(G)$ a set of terminals. We say that a set $S \subseteq V(G)$ is* thin *in G if there is no $v \in S$ such that v*

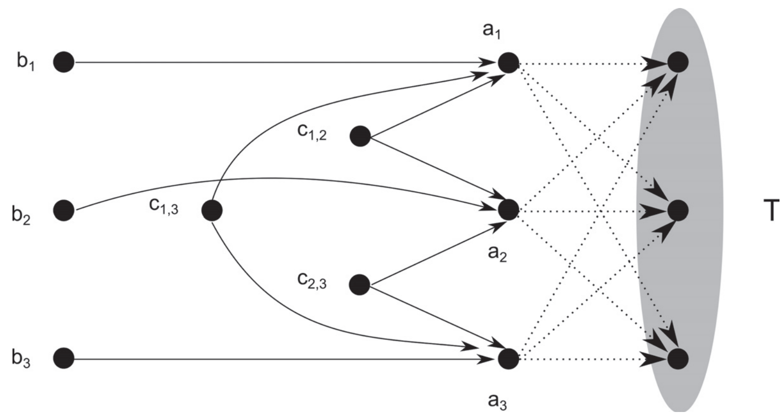

FIG. 5. *An illustration of Remark* 4.8 *in the special case when* $k = 3 = r$.

belongs to the reverse shadow of $S \setminus v$ *with respect to* T.

Refer to Figure 4. The set S is thin because for every $1 \leq i \leq 3$ the vertex w_i does not belong to the reverse shadow of the set $S \setminus \{w_i\}$. However the set $S \cup S'$ is not thin since $(S \cup S') \setminus \{w_1\}$ is a $w_1 - T$ separator, and hence w_1 belongs to the reverse shadow of $(S \cup S') \setminus \{w_1\}$.

THEOREM 4.10 (random sampling). *There is an algorithm* RANDOMSET(G, T, p) *that produces a random set* $Z \subseteq V(G) \setminus T$ *in time* $O^*(4^p)$ *such that the following holds. Let* S *be a* thin *set with* $|S| \leq p$, *and let* Y *be a set such that for every* $v \in Y$ *there is an important* $v - T$ *separator* $S' \subseteq S$. *For every such pair* (S, Y), *the probability that the following two events both occur is at least* $2^{-2^{O(p)}}$:

1. $S \cap Z = \emptyset$, *and*
2. $Y \subseteq Z$.

ALGORITHM 2. RANDOMSET(G, T, p).

1: Enumerate every member of $\mathcal{I}_p$. {See Remark 4.4.}
2: Let $\mathcal{X}$ be the set of exact reverse shadows of members of $\mathcal{I}_p$.
3: Take a random $\mathcal{X}' \subseteq X$ by choosing each element with probability $\frac{1}{2}$, independently at random.
4: Let Z be the union of the exact reverse shadows in $\mathcal{X}'$.
5: **return** Z

Proof. We claim that Algorithm 2 for RANDOMSET(G, T, p) satisfies the requirements. The algorithm RANDOMSET(G, T, p) first enumerates the collection $\mathcal{I}_p$; let $\mathcal{X}$ be the set of all exact reverse shadows of these sets. By Remark 4.4, the size of $\mathcal{X}$ is $O^*(4^p)$, and it can be constructed in time $O^*(4^p)$. Now we show that the set Z satisfies the requirement of the theorem.

Fix a pair (S, Y) as in the statement of the theorem. Let $X_1, X_2, \ldots, X_d \in \mathcal{X}$ be the exact reverse shadows of every member of $\mathcal{I}_p$ that is a subset of S. As $|S| \leq p$, we have $d \leq 2^p$. By assumption that S is *thin*, we have $X_j \cap S = \emptyset$ for every $j \in [d]$. Now consider the following events:

(E1) $Z \cap S = \emptyset$.
(E2) $X_j \subseteq Z$ for every $j \in [d]$.

Note that (E2) implies that $Y \subseteq Z$. Our goal is to show that both events (E1) and (E2) occur with probability $2^{-2^{O(p)}}$.

Let $A = \{X_1, X_2, \ldots, X_d\}$ and $B = \{X \in \mathcal{X} \mid X \cap S \neq \emptyset\}$. By Lemma 4.6, each vertex of S is contained in the exact reverse shadow of at most 4^p members of $\mathcal{I}_p$. Thus $|B| \leq |S| \cdot 4^p \leq p \cdot 4^p$. If no exact reverse shadow from B is selected, then event (E1) holds. If every exact reverse shadow from A is selected, then event (E2) holds. Thus the probability that both (E1) and (E2) occur is bounded from below by the probability of the event that every element from A is selected and no element from B is selected. Note that A and B are disjoint: A contains only sets disjoint from S, while B contains only sets intersecting S. Therefore, the two events are independent, and the probability that both events occur is at least

$$\left(\frac{1}{2}\right)^{2^p}\left(1-\frac{1}{2}\right)^{p\cdot 4^p} = 2^{-2^{O(p)}}. \qquad \square$$

4.3. Derandomization. We now derandomize the process of choosing exact reverse shadows in Theorem 4.10 using the technique of *splitters*. An (n, r, r^2)-splitter is a family of functions from $[n] \to [r^2]$ such that for all $M \subseteq [n]$ with $|M| = r$, at least one of the functions in the family is injective on M. Naor, Schulman, and Srinivasin [21] give an explicit construction of an (n, r, r^2)-splitter of size $O(r^6 \cdot \log r \cdot \log n)$.

THEOREM 4.11 (deterministic sampling). *There exists a randomized algorithm* RANDOMSET(G, T, p) *that produces $t = 2^{2^{O(p)}}$ subsets $Z_1, \ldots, Z_t$ of $V(G) \setminus T$ in time $O^*(2^{2^{O(p)}})$ such that the following holds. Let S be a* thin *set with $|S| \leq p$, and let Y be a set such that for every $v \in Y$ there is an important $v - T$ separator $S' \subseteq S$. For every such pair (S, Y), there is at least one $1 \leq i \leq t$ with*

1. *$S \cap Z_i = \emptyset$, and*
2. *$Y \subseteq Z_i$.*

Proof. In the proof of Theorem 4.10, a random subset of a universe $\mathcal{X}$ of size $n_0 = |\mathcal{X}| \leq 4^p \cdot |V(G)|$ is selected. We argued that for a fixed S, there are a collection $A \subseteq \mathcal{X}$ of $a \leq 2^p$ sets and a collection $B \subseteq \mathcal{X}$ of $b \leq p \cdot 4^p$ sets such that if every set in A is selected and no set in B is selected, then events (E1) and (E2) hold. Instead of selecting a random subset, we construct several subsets such that at least one of them satisfies both (E1) and (E2). Each subset is defined by a pair (h, H), where h is a function in an $(n_0, a+b, (a+b)^2)$-splitter family and H is a subset of $[(a+b)^2]$ of size a (there are $\binom{(a+b)^2}{a} = \binom{(2^p+p4^p)^2}{2^p} = 2^{2^{O(p)}}$ such sets H). For a particular choice of h and H, we select those exact shadows $S \in \mathcal{X}$ into $\mathcal{X}'$ for which $h(S) \in H$. The size of the splitter family is $O((a+b)^6 \cdot \log(a+b) \cdot \log(n_0)) = 2^{O(p)} \cdot \log|V(G)|$, and the number of possibilities for H is $2^{2^{O(p)}}$. Therefore, we construct $2^{2^{O(p)}} \cdot \log|V(G)|$ subsets of $\mathcal{X}$.

By the definition of the splitter, there is a function h that is injective on $A \cup B$, and there is a subset H such that $h(L) \in H$ for every set L in A and $h(M) \notin H$ for every set M in B. For such an h and H, the selection will ensure that (E1) and (E2) hold. Thus at least one of the constructed subsets has the required properties, which is what we wanted to show. $\square$

5. Proof of Lemma 3.7. The goal of this section is to complete the proof of correctness of Algorithm 1 by proving Lemma 3.7. Note that Lemma 3.6 was proved in section 3.

To prove Lemma 3.7, we show that if I is a yes-instance, then there exists a solution S^* for I_1 that remains a solution of the undirected (G_3^*, T, p) as well with probability at least $2^{-2^{O(p)}}$. Suppose that for some solution S^*, the following two properties hold:

1. $Z \cap S^* = \emptyset$, and
2. $r_{G_1,T}(S^*) \bigcup f_{G_1,T}(S^*) \subseteq Z$.

Then Lemma 3.5(2) implies that S^* is a shadowless solution of $I/Z = (G_3, T, p)$. It follows by Lemma 3.1 that S^* is a solution of the undirected instance (G_3^*, T, p) as well. Thus our goal is to prove the existence of a solution S^* for which we can give a lower bound on the probability that these two events occur.

For choosing S^*, we need the following definition.

DEFINITION 5.1 (shadow-maximal solution). *Let (G, T, p) be a given instance of* DIRECTED MULTIWAY CUT. *An inclusionwise minimal solution S is called* shadow-maximal *if $r_{G,T}(S) \bigcup f_{G,T}(S) \bigcup S$ is inclusionwise maximal among all minimal solutions.*

For the rest of the proof, let us fix S^* to be a shadow-maximal solution of instance $I_1 = (G_1, T, p)$ such that $|r_{G_1,T}(S^*)|$ is maximum possible among all shadow-maximal solutions. We now give a lower bound on the probability that $Z \cap S^* = \emptyset$ and $r_{G_1,T}(S^*) \bigcup f_{G_1,T}(S^*) \subseteq Z$. More precisely, we give a lower bound on the probability that all of the following four events occur:

1. $Z_1 \cap S^* = \emptyset$,
2. $r_{G_1,T}(S^*) \subseteq Z_1$,
3. $Z_2 \cap S^* = \emptyset$, and
4. $f_{G_1,T}(S^*) \subseteq Z_2$.

That is, the first random selection takes care of the reverse shadow, the second takes care of the forward shadow, and neither of Z_1 or Z_2 hits S^*. Note that it is somewhat counterintuitive that we choose an S^* for which the shadow is large: intuitively, it seems that the larger the shadow is, the less likely that it is fully covered by Z. However, we need this maximality property in order to give a lower bound on the probability that $Z \cap S^* = \emptyset$.

We want to invoke Theorem 4.10 to obtain a lower bound on the probability that Z_1 contains $Y = r_{G_1,T}(S^*)$ and $Z_1 \cap S^* = \emptyset$. First, we need to ensure that S^* is a *thin* set, but this follows easily from the fact that S^* is a minimal solution.

LEMMA 5.2. *If S is a minimal solution for a* DIRECTED MULTIWAY CUT *instance (G, T, p), then no $v \in S$ is in the reverse shadow of some $S' \subseteq S \setminus \{v\}$.*

Proof. We claim that $S \setminus \{v\}$ is also a solution, contradicting the minimality of S. Suppose that there is a path P from $t_1 \in T$ to $t_2 \in T$, $t_1 \neq t_2$, that intersects S only in v. Consider the subpath of P from v to t_2. As v is in $r(S')$, the set S' is a $v - T$ separator. Thus P goes through $S' \subseteq S \setminus \{v\}$, a contradiction. ∎

More importantly, if we want to use Theorem 4.10 with $Y = r_{G_1,T}(S^*)$, then we have to make sure that for every vertex v of $r_{G_1,T}(S^*)$, there is an important $v - T$ separator that is a subset of S^*. The "pushing argument" of Lemma 5.3 shows that if this is not true for some v, then we can modify the solution in a way that increases the size of the reverse shadow. The choice of S^* ensures that no such modification is possible; thus S^* contains an important separator for every v.

LEMMA 5.3 (pushing). *Let S be a solution of a* DIRECTED MULTIWAY CUT *instance (G, T, p). For every $v \in r(S)$, either there is an $S_v \subseteq S$ which is an important $v - T$ separator, or there is a solution S' such that*

1. $|S'| \leq |S|$,
2. $r(S) \subset r(S')$,
3. $(r(S) \bigcup f(S) \bigcup S) \subseteq (r(S') \bigcup f(S') \bigcup S')$.

Proof. Let $S_0 \subseteq S$ be the subset of S reachable from v without going through any other vertices of S. Then S_0 is clearly a $v-T$ separator. Let S_v be the minimal $v-T$ separator contained in S_0. If S_v is an important $v-T$ separator, then we are done as S itself contains S_v. Otherwise, there exists an important $v-T$ separator S'_v, i.e., $|S'_v| \leq |S_v|$, and $R^+_{G\setminus S_v}(v) \subset R^+_{G\setminus S'_v}(v)$. Now we show that $S' = (S \setminus S_v) \bigcup S'_v$ is a solution for the multiway cut instance. Note that $S'_v \subseteq S'$ and $|S'| \leq |S|$.

First we claim that $r(S) \bigcup (S \setminus S') \subseteq r(S')$. Suppose that there is a path P from β to T in $G \setminus S'$ for some $\beta \in r(S) \bigcup (S \setminus S')$. If $\beta \in r(S)$, then path P has to go through a vertex $\beta' \in S$. As β' is not in S', it has to be in $S \setminus S'$. Therefore, by replacing β with β', we can assume in the following that $\beta \in S \setminus S' \subseteq S_v \setminus S'_v$. By minimality of S_v, every vertex of $S_v \subseteq S_0$ has an incoming edge from some vertex in $R^+_{G\setminus S}(v)$. This means that there is a vertex $\alpha \in R^+_{G\setminus S}(v)$ such that $(\alpha, \beta) \in E(G)$. Since $R^+_{G\setminus S}(v) \subseteq R^+_{G\setminus S'}(v)$, we have $\alpha \in R^+_{G\setminus S'}(v)$, implying that there is a $v \to \alpha$ path in $G \setminus S'$. The edge $\alpha \to \beta$ also survives in $G \setminus S'$ as $\alpha \in R^+_{G\setminus S'}(v)$ and $\beta \in S_v \setminus S'_v$. By assumption, we have a path in $G \setminus S'$ from β to some $t \in T$. Concatenating the three paths, we obtain a $v \to t$ path in $G \setminus S'$ which contradicts the fact that S' contains an (important) $v-T$ separator S'_v. Since $S \neq S'$ and $|S| = |S'|$, the set $S_v \setminus S'_v$ is nonempty. Thus $r(S) \subset r(S')$ follows from the claim $r(S) \bigcup (S \setminus S') \subseteq r(S')$.

Suppose now that S' is not a solution for the multiway cut instance. Then there is a $t_1 \to t_2$ path P in $G \setminus S'$ for some $t_1, t_2 \in T$, $t_1 \neq t_2$. As S is a solution for the multiway cut instance, P must pass through a vertex $\beta \in S \setminus S' \subseteq r(S')$ (by the claim in the previous paragraph), a contradiction. Thus S' is also a minimum solution.

Finally, we show that $r(S) \bigcup f(S) \bigcup S \subseteq r(S') \bigcup f(S') \bigcup S'$. We know that $r(S) \bigcup (S \setminus S') \subseteq r(S')$. Thus it is sufficient to consider a vertex $v \in f(S) \setminus r(S)$. Suppose that $v \notin f(S')$ and $v \notin r(S')$: there are paths P_1 and P_2 in $G \setminus S'$, going from T to v and from v to T, respectively. As $v \in f(S)$, path P_1 intersects S, i.e., it goes through a vertex of $S \setminus S' \subseteq r(S')$; let β be the last such vertex on P_1. Now concatenating the subpath of P_1 from β to v and the path P_2 gives a path from $\beta \in r(S')$ to T in $G \setminus S'$, a contradiction. ☐

Note that if S is a shadow-maximal solution, then solution S' in Lemma 5.3 is also shadow-maximal. Therefore, by the choice of S^*, applying Lemma 5.3 on S^* cannot produce a shadow-maximal solution S' with $r_{G_1,T}(S^*) \subset r_{G_1,T}(S')$, and hence S^* contains an important $v-T$ separator for every $v \in r_{G_1,T}(S)$. Thus by Theorem 4.10 for $Y = r_{G_1,T}(S^*)$, we get the following lemma.

LEMMA 5.4. *With probability at least* $2^{-2^{O(p)}}$, *both* $r_{G_1,T}(S^*) \subseteq Z_1$ *and* $Z_1 \cap S^* = \emptyset$ *occur.*

In the following, we assume that the events in Lemma 5.4 occur. Our next goal is to give a lower bound on the probability that Z_2 contains $f_{G_1,T}(S^*)$. Note that S^* is a solution also of the instance (G_2, T, p): the vertices in S^* remain finite (as $Z_1 \cap S^* = \emptyset$ by the assumptions of Lemma 5.4), and reversing the orientation of the edges does not change the fact that S^* is a solution. Solution S^* is a shadow-maximal solution also in (G_2, T, p): Definition 5.1 is insensitive to reversing the orientation of the edges, and making some of the weights infinite can only decrease the set of potential solutions. Furthermore, the forward shadow of S^* in G_2 is the same as the reverse shadow of S^* in G_1; that is, $f_{G_2,T}(S^*) = r_{G_1,T}(S^*)$. Therefore, assuming that the events in Lemma 5.4 occur, every vertex of $f_{G_2,T}(S^*)$ has infinite weight in G_2. Now we show that S^* contains an important $v-T$ separator in G_2 for every $v \in r_{G_2,T}(S^*) = f_{G_1,T}(S^*)$.

LEMMA 5.5. *If S is a shadow-maximal solution for a* DIRECTED MULTIWAY CUT *instance (G, T, p) and every vertex of $f(S)$ is infinite, then S contains an important $v - T$ separator for every $v \in r(S)$.*

Proof. Suppose to the contrary that there exists $v \in r(S)$ such that S does not contain an important $v - T$ separator. Then by Lemma 5.3, there is another shadow-maximal solution S'. As S is shadow-maximal, it follows that $r(S) \bigcup f(S) \bigcup S = r(S') \bigcup f(S') \bigcup S'$. Therefore, the nonempty set $S' \setminus S$ is fully contained in $r(S) \bigcup f(S) \bigcup S$. However it cannot contain any vertex of $f(S)$ (as they are infinite by assumption) and cannot contain any vertex of $r(S)$ (as $r(S) \subset r(S')$), which is a contradiction. ∎

Recall that S^* is a shadow-maximal solution also in (G_2, T, p). In particular, S^* is a minimal solution for G_2 and so by Lemma 5.2 we have that S^* is thin in G_2 also. Thus Theorem 4.10 can be used (with $Y = r_{G_2,T}(S^*)$) to obtain a lower bound on the probability that $r_{G_2,T}(S^*) \subseteq Z_2$ and $Z_2 \cap S^* = \emptyset$. As the reverse shadow $r_{G_2,T}(S^*)$ in G_2 is the same as the forward shadow $f_{G_1,T}(S^*)$ in G_1, we can state the following lemma.

LEMMA 5.6. *Assuming the events in Lemma* 5.4 *occur, with probability at least $2^{-2^{O(p)}}$ both $f_{G_1,T}(S^*) \subseteq Z_2$ and $Z_2 \cap S^* = \emptyset$ occur.*

Therefore, Lemmas 5.4 and 5.6 imply that with probability at least $(2^{-2^{O(p)}})^2$, the set $Z_1 \bigcup Z_2$ contains $f_{G_1,T}(S^*) \bigcup r_{G_1,T}(S^*)$, and it is disjoint from S^*. Lemma 3.5(2) implies that S^* is a shadowless solution of $I/(Z_1 \bigcup Z_2)$. It follows from Lemma 3.1 that S^* is a solution of the undirected instance (G_3^*, T, p).

LEMMA 5.7. *With probability at least $2^{-2^{O(p)}}$, S^* is a shadowless solution of (G_3, T, p) and a solution of the undirected instance (G_3^*, T, p).*

In summary, with probability at least $2^{-2^{O(p)}}$ Algorithm 1 returns a set S which is a solution of I by Lemma 3.6. This completes the proof of Lemma 3.7.

Appendix A. Bound on the number of important separators (proof of Lemma 4.2). For the proof of Lemma 4.2, first we need to establish some simple properties of important separators, which will allow us to use recursion.

LEMMA A.1. *Let G be a directed graph, and let S be an important $X - Y$ separator. Then the following hold:*

1. *For every $v \in S$, the set $S \setminus v$ is an important $X - Y$ separator in the graph $G \setminus v$.*
2. *If S is an $X' - Y$ separator for some $X' \supset X$, then S is also an important $X' - Y$ separator.*

Proof.

1. Suppose $S \setminus v$ is not a minimal $X - Y$ separator in $G \setminus v$. Let $S_0 \subset S \setminus v$ be an $X - Y$ separator in $G \setminus v$. Then $S_0 \cup v$ is an $X - Y$ separator in G, but $S_0 \cup v \subset S$ holds, which contradicts the fact that S is a minimal $X - Y$ separator in G. Now suppose that there exists an $S' \subseteq V(G) \setminus v$ such that $|S'| \leq |S \setminus v| = |S| - 1$ and $R^+_{(G \setminus v) \setminus (S \setminus v)}(X) \subset R^+_{(G \setminus v) \setminus S'}(X)$. Noting that $(G \setminus v) \setminus (S \setminus v) = G \setminus S$ and $(G \setminus v) \setminus S' = G \setminus (S' \cup v)$, we get $R^+_{G \setminus S}(X) \subset R^+_{G \setminus (S' \cup v)}(X)$. As $|S' \cup v| = |S'| + 1 \leq |S|$, this contradicts the fact that S is an important $X - Y$ separator.
2. As S is an inclusionwise minimal $X - Y$ separator, it is an inclusionwise minimal $X' - Y$ separator as well. Let S' be a witness that S is not an important $X' - Y$ separator in G; i.e., S' is an $X' - Y$ separator such that $|S'| \leq |S|$ and $R^+_{G \setminus S}(X') \subset R^+_{G \setminus S'}(X')$. We claim first that $R^+_{G \setminus S}(X) \subseteq$

$R^+_{G\setminus S'}(X)$. Indeed, if P is any path from X and fully contained in $R^+_{G\setminus S}(X)$, then P is disjoint from S'; otherwise vertices of $P \cap S'$ are in $R^+_{G\setminus S}(X')$, but not in $R^+_{G\setminus S'}(X')$, a contradiction. Next we show that the inclusion $R^+_{G\setminus S}(X) \subset R^+_{G\setminus S'}(X)$ is proper, contradicting that S is an important $X-Y$ separator. As $|S'| \le |S|$, there is a vertex $v \in S \setminus S'$. Since S is a minimal $X-Y$ separator, it has an in-neighbor $u \in R^+_{G\setminus S}(X) \subseteq R^+_{G\setminus S'}(X)$. Now $v \in S$ and $v \notin S'$ imply that $v \in R^+_{G\setminus S'}(X) \setminus R^+_{G\setminus S}(X)$, a contradiction. ∎

Next we show that the size of the out-neighborhood of a vertex set is a submodular function. Recall that a function $f : 2^U \to \mathbb{N} \cup \{0\}$ is *submodular* if for all $A, B \subseteq U$ we have $f(A) + f(B) \ge f(A \cup B) + f(A \cap B)$.

LEMMA A.2 (submodularity). *The function $\gamma(A) = |N^+(A)|$ is submodular.*

Proof. Let $L = \gamma(A) + \gamma(B)$ and $R = \gamma(A \cup B) + \gamma(A \cap B)$. To prove $L \ge R$ we show that for each vertex $x \in V$ its contribution to L is at least as much as its contribution to R. Suppose that the weight of x is w (in our setting, w is either 1 or ∞, but submodularity holds even if the weights are arbitrary). The contribution of x to L or R is either 0, w, or $2w$. We have the following four cases:

1. $x \notin N^+(A)$ and $x \notin N^+(B)$.
   In this case, x contributes 0 to L. It contributes 0 to R as well: every vertex in $N^+(A \cap B)$ or in $N^+(A \cup B)$ is either in $N^+(A)$ or in $N^+(B)$.
2. $x \in N^+(A)$ and $x \notin N^+(B)$.
   In this case, x contributes w to L. To see that x does not contribute $2w$ to R, suppose that $x \in N^+(A \cup B)$ holds. This implies $x \notin A \cup B$, and therefore $x \in N^+(A \cap B)$ can be true only if $x \in N^+(A)$ and $x \in N^+(B)$, which is a contradiction. Therefore, x contributes only w to R.
3. $x \notin N^+(A)$ and $x \in N^+(B)$.
   This is symmetric to the previous case.
4. $x \in N^+(A)$ and $x \in N^+(B)$.
   In this case, x contributes $2w$ to L, and can in any case contribute at most $2w$ to R.

In all four cases the contribution of x to L is always greater than or equal to its contribution to R, and hence $L \ge R$; i.e., γ is submodular. ∎

Recall that $R^+_{G\setminus S}(X)$ is the set of vertices reachable from X in $G\setminus S$. The following claim will be useful for the use of submodularity.

LEMMA A.3. *Let G be a directed graph. If S_1, S_2 are $X-Y$ separators, then both of the sets $N^+(R^+_{G\setminus S_1}(X) \bigcup R^+_{G\setminus S_2}(X))$ and $N^+(R^+_{G\setminus S_1}(X) \bigcap R^+_{G\setminus S_2}(X))$ are also $X-Y$ separators.*

Proof. 1. Let $R_\cap = R^+_{G\setminus S_1}(X) \bigcap R^+_{G\setminus S_2}(X)$ and $S_\cap = N^+(R_\cap)$. As S_1 and S_2 are disjoint from X and Y by definition, we have that $X \subseteq R_\cap$ and Y is disjoint from $R_\cap$. Therefore, every path P from X to Y has a vertex $u \in R_\cap$ followed by a vertex $v \notin R_\cap$, and therefore $v \in S_\cap$. As this holds for every path P, the set $S_\cap$ is an $X-Y$ separator.

2. The argument is the same with the sets $R_\cup = R^+_{G\setminus S_1}(X) \bigcup R^+_{G\setminus S_2}(X)$ and $S_\cup = N^+(R_\cup)$. ∎

Now we prove the well-known fact that there is a unique minimum size separator whose "reach" is inclusionwise maximal.

LEMMA A.4. *There is a unique $X-Y$ separator S^* of minimum size such that $R^+_{G\setminus S^*}(X)$ is inclusionwise maximal.*

Proof. Let λ be the size of a smallest $X-Y$ separator. Suppose to the contrary that there are two separators S_1 and S_2 of size λ such that $R^+_{G\setminus S_1}(X)$ and $R^+_{G\setminus S_2}(X)$ are incomparable and inclusionwise maximal. Let $R_1 = R^+_{G\setminus S_1}(X)$, $R_2 = R^+_{G\setminus S_2}(X)$, $R_\cap = R_1 \cap R_2$, and $R_\cup = R_1 \cup R_2$. By Lemma A.2, γ is submodular and hence

$$\gamma(R_1) + \gamma(R_2) \geq \gamma(R_\cup) + \gamma(R_\cap). \tag{1}$$

As $N^+(R_1) \subseteq S_1$ and $N^+(R_2) \subseteq S_2$, the left-hand side is at most 2λ (in fact, as S_1 and S_2 are minimal $X - Y$ separators, it can be seen that the left-hand side is exactly 2λ). By Lemma A.3, both the sets $N^+(R_\cap)$ and $N^+(R_\cup)$ are $X - Y$ separators. Therefore, the right-hand side is at least 2λ. This implies that equality holds in (1) and in particular $|N^+(R_\cup)| = \lambda$; i.e., $N^+(R_\cup)$ is also a minimum $X - Y$ separator. As $R_1, R_2 \subseteq R_\cup$, every vertex of R_1 and every vertex of R_2 are reachable from X in $G \setminus N^+(R_\cup)$. This contradicts the inclusionwise maximality of the reach of S_1 and S_2. ☐

Let S^* be the unique $X - Y$ separator of minimum size given by Lemma A.4. The following lemma shows that every important $X - Y$ separator S is "behind" this separator S^*.

LEMMA A.5. *Let S^* be the unique $X - Y$ separator of minimum size given by Lemma* A.4. *For every important $X-Y$ separator S, we have $R^+_{G\setminus S^*}(X) \subseteq R^+_{G\setminus S}(X)$.*

Proof. Note that the condition trivially holds for $S = S^*$. Lemma A.4 implies that the only important $X - Y$ separator of minimum size is S^*.

Suppose there is an important $X - Y$ separator $S \neq S^*$ such that $R^+_{G\setminus S^*}(X) \nsubseteq R^+_{G\setminus S}(X)$. Let $R = R^+_{G\setminus S}(X)$, $R^* = R^+_{G\setminus S^*}(X)$, $R_\cap = R \cap R^*$, and $R_\cup = R \cup R^*$. By Lemma A.2, γ is submodular and hence

$$\gamma(R^*) + \gamma(R) \geq \gamma(R_\cup) + \gamma(R_\cap). \tag{2}$$

As $N^+(R^*) \subseteq S^*$, we have that the first term on the left-hand side is at most $|S^*| = \lambda$. By Lemma A.3, the set $N^+(R_\cap)$ is an $X - Y$ separator; hence the second term on the right-hand side is at least λ. It follows that $|N^+(R_\cup)| \leq |N^+(R)| \leq |S|$. Since $R^* \nsubseteq R$ by assumption, we have $R \subset R_\cup$. By Lemma A.3, $N^+(R_\cup)$ is also an $X - Y$ separator, and we have seen that it has size at most $|S|$. Furthermore, $R \subset R_\cup$ implies that any vertex reachable from X in $G \setminus S$ is reachable in $G \setminus N^+(R_\cup)$ as well, contradicting the assumption that S is an important separator. ☐

Now we finally have all the required tools to prove Lemma 4.2.

Proof of Lemma 4.2. Let λ be the size of a smallest $X - Y$ separator. To prove Lemma 4.2, we show by induction on $2p - \lambda$ that the number of important $X - Y$ separators of size at most p is upper bounded by $2^{2p-\lambda}$. Note that if $2p - \lambda < 0$, then $\lambda > 2p \geq p$ and so there is no (important) $X - Y$ separator of size at most p. If $2p - \lambda = 0$, then $\lambda = 2p$. Now if $p = 0$, then $\lambda = p = 0$, and the empty set is the unique important $X - Y$ separator of size at most p. If $p > 0$, then $\lambda = 2p > p$, and hence there is no important $X - Y$ separator of size at most p. Thus we have checked the base case for induction. From now on, the induction hypothesis states that if $X', Y' \subseteq V(G)$ are disjoint sets such that λ' is the size of a smallest $X' - Y'$ separator and p' is an integer such that $(2p' - \lambda') < (2p - \lambda)$, then the number of important $X' - Y'$ separators of size at most p' is upper bounded by $2^{2p'-\lambda'}$.

Let S^* be the unique $X - Y$ separator of minimum size given by Lemma A.4. Consider an arbitrary vertex $v \in S^*$. Note that $\lambda > 0$, and so S^* is not empty. Any important $X - Y$ separator S of size at most p either contains v or not. If S contains v, then by Lemma A.1(1), the set $S \setminus \{v\}$ is an important $X - Y$ separator in $G \setminus v$ of size

Algorithm 3. ImpSep(G, X, Y, p).

Input: A directed graph G, disjoint sets $X, Y \subseteq V$, and an integer p.
Output: A collection of $X - Y$ separators that is a superset of all important $X - Y$ separators of size at most p in G.

Find the minimum $X - Y$ separator S^* of Lemma A.4.
Let $\lambda = |S'|$.
if $p < \lambda$ **then**
  return $\emptyset$
else
  Pick any arbitrary vertex $v \in S^*$.
  Let $\mathcal{S}_1 = \textsc{ImpSep}(G \setminus \{v\}, X, Y, p-1)$.
  Let $\mathcal{S}_1' = \{v \cup S \mid S \in \mathcal{S}_1\}$.
  Let $X' = R^+_{G \setminus S^*}(X) \cup \{v\}$.
  Let $\mathcal{S}_2 = \textsc{ImpSep}(G, X', Y, p)$.
  return $\mathcal{S}_1' \cup \mathcal{S}_2$

at most $p' := p - 1$. As $v \notin X \cup Y \cup V^\infty$, the size λ' of the minimum $X - Y$ separator in $G \setminus v$ is at least $\lambda - 1$. Therefore, $2p' - \lambda' = 2(p-1) - \lambda' = 2p - (\lambda' + 2) < 2p - \lambda$. The induction hypothesis implies that there are at most $2^{2p'-\lambda'} \leq 2^{2p-\lambda-1}$ important $X - Y$ separators of size p' in $G \setminus v$. Hence there are at most $2^{2p-\lambda-1}$ important $X - Y$ separators of size at most p in G that contain v.

Now we give an upper bound on the number of important $X - Y$ separators *not* containing v. By minimality of S^*, vertex v has an in-neighbor in $R^+_{G\setminus S^*}(X)$. For every important $X - Y$ separator S, Lemma A.5 implies $R^+_{G\setminus S^*}(X) \subseteq R^+_{G\setminus S}(X)$. As $v \notin S$ and v has an in-neighbor in $R^+_{G\setminus S^*}(X)$, even $R^+_{G\setminus S^*}(X) \bigcup \{v\} \subseteq R^+_{G\setminus S}(X)$ holds. Therefore, setting $X' = R^+_{G\setminus S^*}(X) \bigcup \{v\}$, the set S is also an $X' - Y$ separator. Now Lemma A.1(2) implies that S is in fact an important $X' - Y$ separator. Since S is an $X - Y$ separator, we have $|S| \geq \lambda$. We claim that in fact $|S| > \lambda$: otherwise $|S| = |S^*| = \lambda$ and $R^+_{G\setminus S^*}(X) \bigcup \{v\} \subseteq R^+_{G\setminus S}(X)$, contradicting the fact that S^* is an important $X - Y$ separator. So the minimum size λ' of an $X' - Y$ separator in G is at least $\lambda + 1$. By the induction hypothesis, the number of important $X' - Y$ separators of size at most p in G is at most $2^{2p-\lambda'} \leq 2^{2p-\lambda-1}$. Hence there are at most $2^{2p-\lambda-1}$ important $X - Y$ separators of size at most p in G that do not contain v.

Adding the bounds in the two cases, we get the required upper bound of $2^{2p-\lambda}$. An algorithm for enumerating all of the at most 4^p important separators follows from the above proof. First, we can find a maximum $X - Y$ flow in time $O(p(|V(G)| + |E(G)|))$ using at most p rounds of the Ford–Fulkerson algorithm, where n and m are the number of vertices and edges of G. It is well known that the separator S^* of Lemma A.4 can be deduced from the maximum flow in linear time by finding those vertices from which Y cannot be reached in the residual graph [12]. Pick any arbitrary vertex $v \in S^*$. Then we branch on whether vertex $v \in S^*$ is in the important separator or not, and we recursively find all possible important separators for both cases. The formal description is given in Algorithm 3. Note that this algorithm enumerates a superset of all important separators: by our analysis above, every important separator appears in either S_1' or S_2, but there is no guarantee that all of the separators in these sets are important. Therefore, the algorithm has to be followed by a filtering phase where we check for each returned separator whether or not it is important. Observe that S is an important $X - Y$ separator if and only if S is the unique minimum $R^+_{G\setminus S}(X) - Y$ separator. As the size of S is at most p, this can be checked in time $O(p(|V(G)| + |E(G)|))$ by finding a maximum flow and constructing the residual graph. The search tree has at most 4^p leaves, and the work to be done in each node is $O(p(|V(G)|+|E(G)|))$. Therefore, the total running time of the branching algorithms is $O(4^p \cdot p(|V(G)| + |E(G)|))$ and returns at most 4^p separators. This is followed by the filtering phase, which takes time $O(4^p \cdot p(|V(G)| + |E(G)|))$. ∎

REFERENCES

[1] N. Bousquet, J. Daligault, and S. Thomassé, *Multicut is* FPT, in Proceedings of the 43rd ACM Symposium on Theory of Computing, 2011, pp. 459–468.

[2] J. Chen, Y. Liu, and S. Lu, *An improved parameterized algorithm for the minimum node multiway cut problem*, Algorithmica, 55 (2009), pp. 1–13.

[3] J. Chen, Y. Liu, S. Lu, B. O'Sullivan, and I. Razgon, *A fixed-parameter algorithm for the directed feedback vertex set problem*, J. ACM, 55 (2008), 21.

[4] R. H. Chitnis, M. Cygan, M. T. Hajiaghayi, and D. Marx, *Directed subset feedback vertex set is fixed-parameter tractable*, in Automata, Languages, and Programming (ICALP 2012), Part 1, Springer-Verlag, Berlin, Heidelberg, 2012, pp. 230–241.

[5] R. H. Chitnis, M. Hajiaghayi, and D. Marx, *Fixed-parameter tractability of directed multiway cut parameterized by the size of the cutset*, in Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, 2012, pp. 1713–1725.

[6] M. Cygan, S. Kratsch, M. Pilipczuk, M. Pilipczuk, and M. Wahlström, *Clique cover and graph separation: New incompressibility results*, in Automata, Languages, and Programming (ICALP 2012), Part 1, Springer-Verlag, Berlin, Heidelberg, 2012, pp. 254–265.

[7] M. Cygan, M. Pilipczuk, M. Pilipczuk, and J. O. Wojtaszczyk, *On multiway cut parameterized above lower bounds*, ACM Trans. Comput. Theory, 5 (2013), 3.

[8] E. Dahlhaus, D. S. Johnson, C. H. Papadimitriou, P. D. Seymour, and M. Yannakakis, *The complexity of multiterminal cuts*, SIAM J. Comput., 23 (1994), pp. 864–894.

[9] R. G. Downey and M. R. Fellows, *Parameterized Complexity*, Springer-Verlag, New York, 1999.

[10] J. Flum and M. Grohe, *Parameterized Complexity Theory*, Springer-Verlag, Berlin, 2006.

[11] L. Ford and D. Fulkerson, *Maximal flow through a network*, Canad. J. Math., 8 (1956), pp. 399–404.

[12] L. R. Ford, Jr., and D. R. Fulkerson, *Flows in Networks*, Princeton University Press, Princeton, NJ, 1962.

[13] N. Garg, V. V. Vazirani, and M. Yannakakis, *Multiway cuts in node weighted graphs*, J. Algorithms, 50 (2004), pp. 49–61.

[14] S. Guillemot, FPT *algorithms for path-transversal and cycle-transversal problems*, Discrete Optim., 8 (2011), pp. 61–71.

[15] D. R. Karger, P. N. Klein, C. Stein, M. Thorup, and N. E. Young, *Rounding algorithms for a geometric embedding of minimum multiway cut*, Math. Oper. Res., 29 (2004), pp. 436–461.

[16] S. Kratsch, M. Pilipczuk, M. Pilipczuk, and M. Wahlström, *Fixed-parameter tractability of multicut in directed acyclic graphs*, in Automata, Languages, and Programming (ICALP 2012), Part 1, Springer-Verlag, Berlin, Heidelberg, 2012, pp. 581–593.

[17] D. Lokshtanov and D. Marx, *Clustering with local restrictions*, Inform. and Comput., 222 (2013), pp. 278–292.

[18] D. Marx, *Parameterized graph separation problems*, Theoret. Comput. Sci., 351 (2006), pp. 394–406.

[19] D. Marx and I. Razgon, *Fixed-parameter tractability of multicut parameterized by the size of the cutset*, in Proceedings of the 43rd ACM Symposium on Theory of Computing, 2011, pp. 469–478.

[20] J. Naor and L. Zosin, *A 2-approximation algorithm for the directed multiway cut problem*, in Proceedings of the 38th Annual Symposium on Foundations of Computer Science, IEEE Computer Society, Washington, DC, 1997, pp. 548–553.

[21] M. Naor, L. J. Schulman, and A. Srinivasan, *Splitters and near-optimal derandomization*, in Proceedings of the 36th Annual Symposium on Foundations of Computer Science, IEEE Computer Society, Washington, DC, 1995, pp. 182–191.

[22] R. Niedermeier, *Invitation to Fixed-Parameter Algorithms*, Oxford University Press, Oxford, UK, 2006.
[23] I. Razgon and B. O'Sullivan, *Almost* 2-SAT *is fixed-parameter tractable*, J. Comput. Syst. Sci., 75 (2009), pp. 435–450.
[24] N. Robertson and P. D. Seymour, *Graph minors*. XIII. *The disjoint paths problem*, J. Combin. Theory Ser. B, 63 (1995), pp. 65–110.
[25] N. Robertson and P. D. Seymour, *Graph minors*. XX. *Wagner's conjecture*, J. Combin. Theory Ser. B, 92 (2004), pp. 325–357.
[26] M. Xiao, *Simple and improved parameterized algorithms for multiterminal cuts*, Theory Comput. Syst., 46 (2010), pp. 723–736.