Fault Tolerant Subgraphs with Applications in Kernelization

William Lochet
University of Bergen, Norway, william.lochet@uib.no

Daniel Lokshtanov
University of California, Santa Barbara, USA, daniello@ucsb.edu

Pranabendu Misra
Max Plank Institute, Germany, pmisra@mpi-inf.mpg.de

Saket Saurabh
Institute of Mathematical Sciences, HBNI, India and and IRL 2000 ReLaX
University of Bergen, Norway, saket@imsc.res.in

Roohani Sharma
Institute of Mathematical Sciences, HBNI, India, roohani@imsc.res.in

Meirav Zehavi
Ben-Gurion University of the Negev, Israel, meiravze@bgu.ac.il

Abstract

In the past decade, the design of *fault tolerant data structures* for networks has become a central topic of research. Particular attention has been given to the construction of a subgraph H of a given digraph D with as fewest arcs/vertices as possible such that, after the failure of any set F of at most $k \geq 1$ arcs, testing whether $D - F$ has a certain property $\mathcal{P}$ is equivalent to testing whether $H - F$ has that property. Here, reachability (or, more generally, distance preservation) is the most basic requirement to maintain to ensure that the network functions properly. Given a vertex $s \in V(D)$, Baswana et al. [STOC'16] presented a construction of H with $\mathcal{O}(2^k n)$ arcs in time $\mathcal{O}(2^k nm)$ where $n = |V(D)|$ and $m = |E(D)|$ such that for any vertex $v \in V(D)$: if there exists a path from s to v in $D - F$, then there also exists a path from s to v in $H - F$. Additionally, they gave a tight matching lower bound. While the question of the improvement of the dependency on k arises for special classes of digraphs, an arguably more basic research direction concerns the dependency on n (for reachability between a pair of vertices $s, t \in V(D)$)—which are the largest classes of digraphs where the dependency on n can be made sublinear, logarithmic *or even constant*? Already for the simple classes of directed paths and tournaments, $\Omega(n)$ arcs are mandatory. Nevertheless, we prove that "almost acyclicity" suffices to eliminate the dependency on n *entirely* for a broad class of dense digraphs called *bounded independence digraphs*. Also, the dependence in k is only a polynomial factor for this class of digraphs. In fact, our sparsification procedure extends to preserve parity-based reachability. Additionally, it finds notable applications in Kernelization: we prove that the classic Directed Feedback Arc Set (DFAS) problem as well as Directed Edge Odd Cycle Transversal (DEOCT) (which, in sharp contrast to DFAS, is W[1]-hard on general digraphs) admit polynomial kernels on bounded independence digraphs. In fact, for any $p \in \mathbb{N}$, we can design a polynomial kernel for the problem of hitting all cycles of length ℓ where ($\ell \mod p = 1$). As a complementary result, we prove that DEOCT is NP-hard on tournaments by establishing a combinatorial identity between the minimum size of a feedback arc set and the minimum size of an edge odd cycle transversal. In passing, we also improve upon the running time of the sub-exponential FPT algorithm for DFAS in digraphs of bounded independence number given by Misra et at. [FSTTCS 2018], and give the first sub-exponential FPT algorithm for DEOCT in digraphs of bounded independence number.

2012 ACM Subject Classification Theory of computation → Data structures design and analysis; Theory of computation → Parameterized complexity and exact algorithms; Theory of computation → Sparsification and spanners

Keywords and phrases sparsification, kernelization, fault tolerant subgraphs, directed feedback arc set, directed edge odd cycle transversal, bounded independence number digraphs

Funding *Daniel Lokshtanov*: European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant no. 715744), and United States - Israel Binational Science Foundation grant no. 2018302.
Saket Saurabh: European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant no. 819416), and Swarnajayanti Fellowship grant DST/SJF/MSA-01/2017-18.
Meirav Zehavi: Israel Science Foundation (ISF) individual research grant (grant no. 1176/18) and Israel Binational Science Foundation (BSF) start up grant.

1 Introduction

In most real-life applications, even the most reliable networks are highly prone to unexpected failures of a small number of links that connect their nodes. In the past decade, the design of *fault tolerant data structures* for networks has become a central topic of research [7, 9, 14, 53, 50, 12, 17, 18, 19, 11, 28, 51, 52]. Generally, the scenario under study concerns the design of a structure that, after the failure of any set F of at most $k \geq 1$ arcs (representing links) in a given digraph D (representing a network), should provide a fast answer to certain types of queries that address the properties of $D - F$. The most common queries of this form address the *reachability* between two vertices, or, more generally, the length of a shortest path existent, if any, between them. Indeed, reachability (or, more generally, distance preservation) is the most basic requirement to maintain to ensure that the network functions properly. In this context, particular attention has been given to the case where the data structure should consist of a subgraph or a minor of D with as fewest arcs/vertices as possible [7, 53, 9, 11, 8, 51, 14]. Then, queries can be answered by standard means as the usage of BFS or Dijkstra's algorithm. In particular, these simple data structures are of interest as they also double as *sparsifiers*. The study of various graph sparsifiers—such as *flow-sparsifiers* [42] which are closely related to the aforementioned data structures—is a fundamental, active area of research in computer science and structural graph theory [24, 5, 32, 42, 16].

More concretely, in the Fault-Tolerance (S,T)-Reachability problem (or FTR(S,T) for short), we are given a digraph D, two (not necessarily disjoint) terminals sets $S, T \subseteq V(D)$, and a positive integer k. The objective is to construct a subgraph H of D with minimum number of arcs/vertices such that, after the failure of any set of at most k arcs in D, the following property is preserved for any two vertices $s \in S$ and $t \in T$: if there still exists a directed path from s to t in D, then there also still exists a directed path from s to t in H. Clearly, a trivial lower bound on the number of arcs in H is $m = \Omega(n^2)$. For the case where $|S| = 1$ and $T = V(D)$, Baswana et al. [9] presented a construction of a subgraph H with $\mathcal{O}(2^k n)$ arcs in time $\mathcal{O}(2^k nm)$ where $n = |V(D)|$ and $m = |E(D)|$. Additionally, they gave a tight matching lower bound: for any $n, k \in \mathbb{N}$ where $n \geq 2^k$, there exists a digraph on n vertices where H must have $\Omega(2^k n)$ arcs.

Naturally, the question of the improvement of the dependency on k arises for special classes of digraphs. However, an arguably more radical research direction to pursue concerns the dependency on n.

> Which are the largest classes of digraphs for which FTR(S,T) admits subgraphs whose size dependency on n can be made sublinear, logarithmic *or even constant*?

At first glance, when we consider the simplest sparsest digraph existent, this pursuit seems futile. Indeed, already in the case where $S = \{s\}$, $T = \{t\}$, $k = 1$ and D is a directed path from s to t, the only solution is to choose $H = D$. At second glance, when we consider the simplest densest digraph existent, again we reach a dead-end: for S, T and k as before, define D as the tournament obtained by adding, to a directed path $s = v_1 \to v_2 \to \ldots \to v_n = t$, all arcs going from v_i to v_j for every $j + 1 < i$; then, to construct H, we must select the entire path.

We show that "almost acyclicity" suffices to eliminate the dependency on n *entirely* for a broad class of dense digraphs called *bounded independence number digraphs*. Furthermore, one can achieve a polynomial dependence in terms of k for this digraph class.

To step beyond the strict confinement of tournaments where *all* relations (arcs) between the input entities (vertices) must be both *present* and *known*, Fradkin and Seymour [36] initiated the study of bounded independence digraphs. Formally, for any integer $\alpha \geq 1$, the class of α-bounded independence digraphs, denoted by $\mathcal{D}_\alpha$, is defined as follows.

$$\mathcal{D}_\alpha = \{D \mid D \text{ is a digraph and the maximum size of an independent set in } D \text{ is at most } \alpha\}.$$

For this class of digraphs, Fradkin and Seymour [36] studied the k-Disjoint Paths problem, and showed that it admits a polynomial time algorithm for any fixed value of k. Observe that $\mathcal{D}_\alpha$ is *hereditary*, and for $\alpha = 1$, it coincides with the class of tournaments. Furthermore, even for $\alpha = 2$, it contains digraphs with a linear fraction of vertex pairs that have no arc between them—thus, it can accommodate the lack of a large number of links/relations.

Our main technical contribution is the following combinatorial lemma.

▶ **Lemma 1.1.** *Given a digraph $D \in \mathcal{D}_\alpha$, positive integers k and ℓ, and $S \subseteq V(D)$ such that every strongly connected component of $D - S$ has at most ℓ vertices, the* Fault-Tolerance (S, S)-Reachability (FTR(S, S)) *problem admits a solution H on $|S|^2(k\ell)^{\mathcal{O}(4^{\alpha\ell^2})}$ vertices. Furthermore, such a solution H can be found in polynomial time.*

In particular, when $D - S$ is acyclic, $\ell = 1$. Thus, if $|S|$ and ℓ are independent of n (such as the case where $|S| = |T| = \ell = 1$ discussed earlier), the dependency on n is eliminated. (We remark that a solution for Fault-Tolerance (S, T)-Reachability where $S \neq T$ is subsumed by a solution for Fault-Tolerance $(S \cup T, S \cup T)$-Reachability.) Note that we extend the class of digraphs dealt with beyond acyclicity at two fronts: enabling S to be a modulator, thus $D - S$ rather than D should be "almost acyclic"; enabling the strongly connected components to be of size that is ("small" but) larger than 1.

In fact, our result generalizes to *parity reachability*. More precisely, in the Fault-Tolerance (S, T)-Parity Reachability problem, we are given a digraph D, two terminal sets $S, T \subseteq V(D)$, positive integers k and p, and a non-negative integer r. The objective is to construct a subgraph H of D with as few arcs/vertices as possible, such that, after the failure of any set of at most k arcs in D, the following property is preserved for any two vertices $s \in S$ and $t \in T$: if there exists a directed path from s to t in D whose length q satisfies $(q \mod p = r)$, then there also exists a directed path from s to t in H whose length q' satisfies $(q' \mod p = r)$. For this problem, we prove the following combinatorial lemma.

▶ **Lemma 1.2.** *Given a digraph $D \in \mathcal{D}_\alpha$, positive integers k, ℓ, p, a non-negative integer r, and $S \subseteq V(D)$ such that every strongly connected component of $D - S$ has at most ℓ vertices, the* Fault-Tolerance (S, S)-Parity Reachability *problem admits a solution H on $(|S|\alpha\ell pk)^{\mathcal{O}(4^{\alpha\ell^2})}$ vertices. Furthermore, such a solution H can be found in polynomial time.*

1.1 Applications in Kernelization

Directed Feedback Arc Set. From the perspective of Parameterized Complexity, with the exception of Directed Multicut, the Directed Feedback Arc/Vertex Set (DFA/VS) problem is the most well studied parameterized problem on digraphs. (On general digraphs, the vertex and arc versions of the problem are equivalent [25].) Formally, this problem is defined as follows.

Directed Feedback Arc Set (DFAS)	**Parameter:** k
Input: A digraph D and a non-negative integer k.	
Question: Does there exist $S \subseteq E(D)$ of size at most k such that $D - S$ is a DAG?	

We remark that this problem is among Karp's 21 original NP-complete problems [39]. Already a decade ago, the DFAS problem has been shown to be *fixed-parameter tractable (FPT)* parameterized by the solution size k [20]. Specifically, Chen et al. [20] developed an algorithm that solves DFAS in time $\mathcal{O}(k!4^k k^4 mn)$, based on the powerful machinery of important separators [25]. Since then, the quest to assert the existence of a polynomial kernel for this problem has been unfruitful. Over the years, it has been repeatedly posed as a major challenge in the subfield of Kernelization [25, 30, 47, 46] (also see [1] for a number of workshops and schools where it was posed as an open problem). In fact, the two specific problems whose polynomial kernelization complexity is completely unknown and their resolution is raised most frequently are DFAS and Multiway Cut [25, 30]. At the front of parameterized algorithms, the recent work by Lokshtanov et al. [44] improved upon the polynomial factor of the aforementioned algorithm by the design of an $\mathcal{O}(k!4^k k^5(m+n))$-time algorithm. It is known that unless the Exponential Time Hypothesis (ETH) is false, parameterized by the treewidth tw of the underlying undirected graph, DFAS cannot be solved in time $2^{o(\mathsf{tw}\log \mathsf{tw})} \cdot n^{\mathcal{O}(1)}$. However, it is unknown whether DFAS is solvable in time $2^{o(k \log k)} \cdot n^{\mathcal{O}(1)}$. In this regard, the only lower bound known is of $2^{\Omega(k)} \cdot n^{\mathcal{O}(1)}$ under the ETH [25, 44].

Particular attention has been given to the parameterized complexity of DFAS on tournaments. The classical complexity (NP-hardness) of DFAS on tournaments has a curious history. More than two decades ago, this problem was conjectured to be NP-hard by Bang-Jensen and Thomassen [6]. In 2008, Ailon et al. [2] proved that this problem does not admit a polynomial-time algorithm unless NP$\subseteq$BPP. Later, the reduction of Ailon et al. [2] was derandomized independently by Alon [3] and Charibt et al. [15], to prove that DFAS on tournaments is NP-hard. With respect to Parameterized Complexity, Alon et al. [4] proved that DFAS on tournaments admits a sub-exponential time parameterized algorithm (with running time $2^{\mathcal{O}(\sqrt{k}\log^2 k)} \cdot n^{\mathcal{O}(1)}$), to which end they introduced the method of chromatic coding. Later, the $\log^2 k$ factor in the exponent was shaved in independent works by Feige [33] and Karpinski and Schudy [40]. Fomin and Pilipczuk [35] presented a general approach, based on a bound on the number of k-cuts in transitive tournaments, to achieve the same running time for DFAS on tournaments. Based on this approach, Misra et al. [48] developed a sub-exponential time parameterized algorithm for DFAS on digraphs in $\mathcal{D}_\alpha$, with running time $2^{\mathcal{O}(\alpha^2\sqrt{k}\log(\alpha k))} \cdot n^{\mathcal{O}(\alpha)}$. Yet, the (arguably more) intriguing question of the existence of a polynomial kernel for DFAS on digraphs in $\mathcal{D}_\alpha$ remained unsolved. On tournaments, Bessy et al. [10] have proved that DFAS admits a linear-vertex kernel (improving upon polynomial kernels given in [4, 29]). Based on our combinatorial lemma (Lemma 1.1), we establish the following theorem.

▶ **Theorem 1.3.** DFAS *on* $\mathcal{D}_\alpha$ *admits a kernel of size* $k^{\mathcal{O}(4^\alpha)}$.

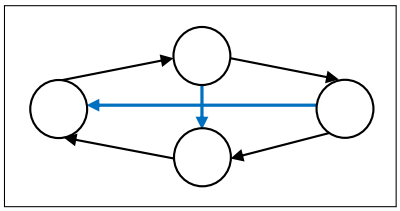

Figure 1 A directed edge odd cycle transversal (in blue) that is not a directed feedback arc set.

In addition to its rich history in theoretical studies, the elimination of directed feedback loops is highly relevant to rank aggregation, Voting Theory, the resolution of inconsistencies in databases, and the prevention of deadlocks [57, 10, 38, 41, 20, 34]. While in a wide-variety of applications, *most* relations between the entities in a network are both present and known, it is generally unrealistic (in real-world partial and noisy data) that *all* relations will be so. Then, the usage of a bounded independence digraphs naturally comes into play. In passing, using Theorem 1.3, we also improve the running time for DFAS on digraphs in $\mathcal{D}_\alpha$, given by Misra et al. [48], by eliminating the dependence of α in the exponent of n. That is, we have the following theorem.

▶ **Theorem 1.4.** DFAS *on* $\mathcal{D}_\alpha$ *can be solved in* $2^{f(\alpha)\sqrt{k}\log k} \cdot n^{\mathcal{O}(1)}$*, where* $f(\alpha)$ *is some function of* α *and* n *is the number of vertices in* D*.*

Directed Edge Odd Cycle Transversal. The Directed Edge Odd Cycle Transversal (DEOCT) problem is the parity-based version of DFAS, formally defined as follows. (On general digraphs, the vertex and arc versions of the problem are equivalent [45]).

Directed Edge Odd Cycle Transversal (DEOCT)	**Parameter:** k
Input: A digraph D and a non-negative integer k.	
Question: Does there exist $S \subseteq E(D)$ of size at most k such that $D-S$ has no odd cycle?	

Observe that a tournament has no directed cycle if and only if it has no directed triangle (a cycle on three vertices). In turn, this simple observation implies that, given a tournament D, any subset S of the *vertices* of D has the following property: $D - S$ is a DAG if and only if it has no directed odd cycle. Thus, the vertex versions of DFAS and DEOCT on tournaments are equivalent. However, for DFAS and DEOCT the situation is not so clear. Indeed, it is not difficult to come up with a tournament D and a subset of *arcs* S of D such that $D - S$ is not a DAG, yet it has no directed odd cycle (see, e.g., Fig. 1). Nonetheless, we are able to prove that given a tournament D and a subset S of the *arcs* of D such that $D - S$ has no directed odd cycle, there exists a subset of arcs S' of D such that $D - S'$ is a DAG and $|S'| \leq |S|$. In particular, we thus establish the following result.

▶ **Theorem 1.5.** DEOCT *on tournaments is* NP*-hard.*

The question of the parameterized complexity of DEOCT was explicitly stated as an open problem [26] for the first time in 2007, immediately after the announcement of the first parameterized algorithm for DFAS. Since then, the problem has been re-stated several times [22, 23, 47, 46]. Recently, Lokshtanov et al. [45] proved that DEOCT is W[1]-hard. Specifically, this means that DEOCT is highly unlikely to be FPT or admit a kernel of any size (even exponential in k). Based on the parity-based generalization of our combinatorial lemma (Lemma 1.2), we establish a polynomial kernel for DEOCT on $\mathcal{D}_\alpha$, which stands in sharp contrast to its aforementioned status on general digraphs.

▶ **Theorem 1.6.** DEOCT *on* $\mathcal{D}_\alpha$ *admits a kernel of size* $(\alpha k)^{\mathcal{O}(4^{4\alpha^3})}$.

In fact, we present combinatorial results stronger than Lemma 1.2 that yield a polynomial kernel for a more general version of DEOCT, where instead of hitting directed odd cycles, the objective is to hit directed cycles whose length ℓ satisfies ($\ell \mod p = 1$) for an integer $p \in \mathbb{N}$ given as input.[1]

Modulo p Directed Cycle Transversal (mod(p)-DCT)	**Parameter:** k
Input: A digraph D and non-negative integers k and p.	
Question: Does there exist $S \subseteq E(D)$ of size at most k such that $D - S$ has no cycle of length $1 \mod p$?	

▶ **Theorem 1.7.** mod(p)-DCT *on* $\mathcal{D}_\alpha$ *admits a kernel of size* $(p\alpha k)^{\mathcal{O}(4^{\alpha^3 p^2})}$.

Having Theorem 1.6 at hand, we also show how to employ the general approach of Fomin and Pilipczuk [35] to derive a sub-exponential time parameterized algorithm for DEOCT on digraphs in $\mathcal{D}_\alpha$.

▶ **Theorem 1.8.** DEOCT *on* $\mathcal{D}_\alpha$ *admits an algorithm with running time* $2^{\mathcal{O}(f(\alpha)\sqrt{k}\log k)} \cdot n^{\mathcal{O}(1)}$, *where* $f(\alpha)$ *is a function of* α *and* n *is the number of vertices in* D.

1.2 Towards the proof of Lemmas 1.1 and 1.2: Cut Preserving Sets and Parity Reserving Sets

The most central notion in this paper is of a *cut preserving set*. Informally, for a digraph D, a pair of vertices s, t and an integer k, a set $\mathcal{Z} \subseteq V(D)$ is called a k-cut preserving set[2] for (s,t) in D if it preserves all (s,t)-arc cuts of size at most k. That is, A is an (s,t)-arc cut with at most k arcs in D if and only if A is a such a cut in $D[\mathcal{Z}]$. Observe that the graph induced on such a k-cut preserving set $\mathcal{Z}$ is a candidate solution for FTR($\{s\}, \{t\}$) problem. Clearly $V(D)$ is a k-cut preserving set for any pair of vertices s, t. The intent is to have such a set of "small" size. Towards this, let us discuss some properties that suffice for $\mathcal{Z}$ to be a k-cut preserving set for (s,t) in D.

Since $\mathcal{Z} \subseteq V(D)$, any (s,t)-arc cut of D is an (s,t)-arc cut of $D[\mathcal{Z}]$. For the other direction, we need the property that, for any $A \subseteq E(D)$ of size at most k, the existence of an (s,t)-path in $D - A$ implies the existence of an (s,t)-path in $D[\mathcal{Z}] - A$. Let us now see which properties suffice to imply the above property. We begin with a special case. Suppose there is a "large" flow from s to t in D. In particular, suppose there are at least $k+1$ internally vertex-disjoint (s,t)-paths in D. Then, in $\mathcal{Z}$ it is enough to keep the vertices of some $k+1$ vertex-disjoint (s,t)-paths, as no arc set of size at most k can hit all these paths. The more involved case occurs when the flow from s to t in D is at most k. Consider any (s,t)-path P in D. Ideally (if we did not have a size constraint on $\mathcal{Z}$) we would have preserved all the vertices of P in $\mathcal{Z}$. Clearly, this can be expensive in terms of the size of $\mathcal{Z}$. Nevertheless, we can merge the ideas above (the "large-flow idea" and the "keep-full-path idea") to get the desired result. To see this, let P be a (s,t)-path in D. Let $\mathcal{Z}$ be a set of vertices such that, either all the vertices of P are in $\mathcal{Z}$ or if the vertices of a (u,v)-subpath of P are not

[1] Note that a fundamental difference between this result and Lemma 1.2 is that the latter only works for any modulo and not just 1. The reason for this is explained in Section 6.

[2] This is not the way it is defined later. However, for the sake of exposition, we start with this definition and refine it to have properties that also guarantee this property implicitly.

in $\mathcal{Z}$, then there are $k+1$ internally vertex-disjoint (u,v)-paths in $D[\mathcal{Z}]$. That is, if the vertices of a subpath are missing in $\mathcal{Z}$, then $\mathcal{Z}$ contains a witness of a large flow for the endpoints of this subpath. Observe that such a set $\mathcal{Z}$ suffices to be a k-cut preserving set for (s,t) in D. This is because if P is an (s,t)-path in $D-A(A \subseteq E(D)$ and $|A| \leq k)$, then either all the vertices of P are in $\mathcal{Z}$ or for any missing (u,v)-subpath of P, since there are $k+1$ vertex-disjoint (u,v)-paths in $D[\mathcal{Z}]$, at least one still remains in $D[\mathcal{Z}]-A$. Thus, in $D[\mathcal{Z}]-A$, one can find an (s,t)-path: for the missing subpaths of P in $\mathcal{Z}$, there exists some (other) path between the same endpoints in $D[\mathcal{Z}]-A$ which together yield an (s,t)-walk (and hence an (s,t)-path) in $D[\mathcal{Z}]-A$. These properties are formalized in Definition 3.1.

1.2.1 About Computing k-Cut Preserving Sets

Next we give an intuition for how one can compute such k-cut preserving sets for a digraph $D \in \mathcal{D}_\alpha$, each of whose strongly connected component has size at most ℓ. For exposition purposes, consider (for now), only the case where D is acyclic (i.e. $\ell = 1$). With a certain technical argument, the general case reduces to this one. Moreover, we use the definition of a k-cut preserving set from the beginning of this section for this illustration as it allows us to convey our ideas in a clearer manner.

The proof will use induction on α. As the base case, consider the case when $\alpha = 1$, that is, D is a transitive tournament. As D is transitive, there exists a topological ordering of the vertices of D. Consider the set S of vertices between s and t in this ordering. Note that any path from s to t only uses vertices in S. So, either S is smaller than $k+1$, and then $S \cup \{s,t\}$ is a k-cut preserving set for (s,t), or it can be seen that there is no arc-cut for (s,t) of size at most k. In the latter case, the union of $\{s,t\}$ and any subset of $k+1$ vertices of S is a k-cut preserving set for (s,t); indeed, in the subgraph induced by the union there is still no arc-cut for (s,t) of size at most k.

Now, let us hint at how the inductive step of the proof works. First, we note that, if $P_1, \ldots, P_{k+1}$ are $k+1$ internally vertex-disjoint (s,t)-paths, then $\mathcal{Z} = \cup_{i \in [k+1]} P_i$ is a k-cut preserving set for the pair (s,t), because there is no arc-cut of (s,t) in both D and $D[\mathcal{Z}]$ of size at most k. Moreover, since D is acyclic and $D \in \mathcal{D}_\alpha$, if these paths exist, then Observation 2.1 implies that we can assume that all these paths are shorter than $2\alpha+1$ and thus $|\mathcal{Z}| \leq k(2\alpha+1)$.

The last argument means that we can assume the existence of a (s,t)-vertex cut of size at most k. For simplicity, suppose that $\{c_1, c_2\}$ is a minimal (s,t)-vertex -cut. Since $\{c_1, c_2\}$ is a vertex cut, any path from s to t in D can be decomposed as a path from s to c_i, a path from c_i to c_j and then a path from c_j to t, where i and j are two indices (possibly equal) in $\{1,2\}$. Here, we mean that none of the three paths contains c_i (or c_j) as an internal vertex. For $i \in \{1,2\}$, let S_i be the union of the set of vertices of the paths from s to c_i that intersect $\{c_1, c_2\}$ only on the last vertex, and T_i be the union of the set of vertices of the paths from c_i to t that intersect $\{c_1, c_2\}$ only on the first vertex. Finally, for distinct $i, j \in \{1,2\}$, let $C_{i,j}$ be the union of the set of vertices of the paths from c_i to c_j. Because of the last remark on how any path from s to t can be decomposed, taking the union of six k-cut preserving sets-namely, for each $i, j \in \{1,2\}$, $i \neq j$, for (s, c_i) in $D[S_i]$, (c_i, t) in $D[T_i]$ and (c_i, c_j) in $D[C_{i,j}]$- gives a k-cut preserving set for (s,t) in D. Now, the question is how to use the induction hypothesis to find a k-cut preserving set for each of these pairs. Consider first the digraph induced by the vertices in S_1. Because $\{c_1, c_2\}$ is a minimal (s,t)-vertex cut, the only vertices of S_1 that can possibly have "outgoing arcs towards" t in S_1 are s and c_1. Moreover, since $\{c_1, c_2\}$ is a minimal (s,t)-vertex cut, there exists a path from c_1 to t in D and thus t is reachable from any vertex of S_1. However, since D is acyclic, this means that

there is no arc from t to any of the vertices of S_1, else we would get a closed walk and thus a cycle. This implies that $D[S_1 \setminus \{s, c_1\}] \in \mathcal{D}_{\alpha-1}$ as any independent set of $S_1 \setminus \{s, c_1\}$ can be extended with t. We cannot apply the induction hypothesis to find a k-cut preserving set for (s, c_1) in S_1 because the independence number of $D[S_1]$ could be equal to α, however the above shows the spirit of the arguments that will be used to find subgraphs with smaller independence number where we can apply the induction hypothesis. A similar argument would also give that the independence number of $D[T_1 \setminus \{c_1, t\}]$ is at most $\alpha - 1$ as any independent set can be extended using s.

The previous argument does not apply to $C_{1,2}$, because the vertices of $C_{1,2}$ can be adjacent to s or t (some vertices of $C_{1,2}$ can be adjacent to s and some can be adjacent to t). This is the case that requires a stronger and more technical definition for a k-cut preserving set. In particular, we need to understand what happens to the vertices of D that are on a path from s to t but do not belong to a k-cut preserving set for this pair.

1.2.2 Preserving Length Modulo p

As explained earlier, our method allows us not only to preserve the existence of a path after the removal of k arcs, but the existence of a path of certain length modulo p. The argument follows the same lines as what we just described, but need the following additional observation. Suppose s and t are two vertices of a digraph $D \in \mathcal{D}_\alpha$ and there exists $p^2\alpha$ vertex-disjoint (s, t)-paths in D. By the pigeonhole principle, $p\alpha$ of those paths must have the same length modulo p. Let $P_1, \ldots, P_{p\alpha}$ denote those paths and X denote the set of vertices appearing just after s along those paths. X is a set of size of size $p\alpha$, and because the largest independent set of D is at most α, it means that the chromatic number of $D[X]$ is at most p. By Gallai-Roy Theorem [37, 56], there exists a path P of length $p-1$ in $D[X]$. Using this path and the P_i, we are able to find a path of *any* length modulo p from s to t (see the proof of Lemma 2.2 for all details). This implies that if s and t are two vertices with more than $p^2\alpha + k$ vertex-disjoint paths from s to t, then preserving exactly $p^2\alpha + k$ of these paths is enough for our purpose. Indeed, after the removal of k arcs, there will still be $p^2\alpha$ vertex-disjoint (s, t) paths and thus a path of every parity. The rest of the argument is identical.

1.3 Deriving Polynomial Kernels for DFAS and DEOCT

Let us now briefly explain how to derive a polynomial kernel for DFAS when the input digraph belongs to $\mathcal{D}_\alpha$, from our result on fault-tolerant subgraphs. First note that if $D \in \mathcal{D}_\alpha$ then every induced cycle in D has length at most $2\alpha + 1$. Let (D, k) be an instance of DFAS, and consider a maximal set of arc disjoint induced cycles in D. If this set consists of more than k cycles, then any solution to (D, k) has to pick one arc per cycle, and (D, k) is a NO instance. If not, let S be the union of these cycles. S is a set of less than $(2\alpha + 1) \cdot k$ vertices such that $D - S$ is acyclic. Therefore, we can apply our result to find a solution H to the problem of Fault-Tolerance (S, S)-Reachability of size at most $|S|^2 k^{\mathcal{O}(4^\alpha)}$. We claim that H is the desired kernel. Indeed, suppose that A is a set of arcs such that $H - A$ is acyclic, but $D - A$ contains a cycle. By construction of S, this cycle must use vertices of S. However, we know that if a path exists between two vertices of S in $D - A$, then such a path also exists in $H - A$. This implies the existence of a closed walk in $H - A$, a contradiction.

Using arguments similar as above, together with our Lemma 1.2, one can design a polynomial kernel for DEOCT.

More related works.

In addition to the results mentioned above, let us mention a few more related works. Fault tolerant data-structures for various graph theory problems, and more generally dynamic graph data-structures, are well-studied. A seminal work of Nagamochi and Ibaraki [49] provides a fault-tolerant reachability subgraph of size $\mathcal{O}(nk)$ for k-edge faults in an undirected graph on n-vertices, for any choice of sources and sinks. As we mentioned earlier, Baswana et.al. [9] give a single source fault tolerant reachability subgraph of size $\mathcal{O}(2^k n)$ for k vertex or edge faults. A simpler algorithm for this problem was presented in [43], using *important separators* [25]. Bodwin and Greg [13] provided a characterization of graph families for which distance preservers for p pairs with $\mathcal{O}(n+p)$ edges is guaranteed to exist. A few results are also known about fault tolerant reachability data-structures. Patrascu and Thorup [54] present a data structure of size $\mathcal{O}(m)$, that can process any set F of k edge faults in a graph G in $\mathcal{O}(k \log^2 n \log\log n)$ time, and subsequently answer reachability queries for pairs of vertices in $G - F$ in time $\mathcal{O}(\log\log n)$. For vertex failures, in recent work Duan and Pettie [31] present a data structure of size $\mathcal{O}(mk \log n)$, which can process any set of k edge faults in $\mathcal{O}(k^3 \log^3 n)$ time, and then answer reachability queries in $\mathcal{O}(k)$ time. They also present an improved randomized algorithm.

Roadmap to the paper

Section 2 contains introduction to some basic terminology and notation, and also some offhand observations and propositions that will be used throughout. Section 3 contains the algorithm to compute k-cut preserving sets. Section 4 contains some applications of the result in Section 3. More specifically, it contains the proof of Lemma 1.1 and Theorem 1.3. Section 5 contains the definitions of k-parity preserving sets together with an algorithm to compute them. Section 6 them applies the result of Section 5 to prove Lemma 1.2 and Theorem 1.7. Section 7 proves Theorem 1.5. Section 8 proves Theorems 1.4 and 1.8.

2 Preliminaries

For standard notations and terminology that is not defined here, we refer to [27].
Sets: For positive integer i, j, $[i]$ denotes the set $\{1, \ldots, i\}$ and $[i, j]$ denote the set $\{i, i+1, \ldots, j\}$. For a set S, S^2 denotes the set of ordered pairs of S, that is $S^2 = \{(u, v) \mid u \in S, v \in S\}$.
Digraphs: For a digraph D, $V(D)$ denotes the vertex set of D and $E(D)$ denotes the arc set of D. For any $X \subseteq V(D)$ (resp. $X \subseteq E(D)$), $D - X$ denotes the digraph obtained by deleting the vertices (resp. edges) of X. For any $v \in V(D)$, $N^+_D(v)$ (resp. $N^-_D(v)$) denotes the set of out-neighbours (resp.in-neighbours) of v in D, that is $N^+_D(v) = \{u \in V(D) \mid (v, u) \in E(D)\}$ (resp. $N^-_D(v) = \{u \in V(D) \mid (u, v) \in E(D)\}$). Whenever the digraph D is clear from the context, we drop the subscript D in $N^+_D(v)$ (resp. $N^-_D(v)$). For any $X, Y \subseteq V(D)$, $E(X, Y)$ denotes the set of arcs of D with tail in X and head in Y, that is, $E(X, Y) = \{(u, v) \in E(D) \mid u \in X, v \in Y\}$. A digraph D is called *strongly connected* if for each $u, v \in V(D)$ there is a path from u to v and, a path from v to u in D. A set $X \subseteq V(D)$ is called a *strongly connected component* of D if $D[X]$ is a strongly connected digraph and for each $X' \supseteq X$, $D[X']$ is not a strongly connected digraph. A *tournament* is a digraph where there is exactly one arc between each pair of vertices. A digraph with no cycles is called a *directed acyclic graph (dag)*. A tournament with no cycles is called a *transitive tournament*.
Paths: A *path* P is a graph such that there exists an ordering $(v_1, \ldots, v_q)$ of its vertex set

$V(P)$ such that $E(P) = \{(v_i, v_{i+1}) \mid i \in [q-1]\}$. Such a path is called a (v_1, v_q)-*path*, v_1, v_q are called the *end-points of* P and $v_2, \ldots, v_{q-1}$ are called the *internal vertices of* P. A path P is *even* (resp. odd) if the number of arcs/edges in it is even (resp. odd). We say that P is a path in the digraph D if P is a subgraph of D. We say that P is an *induced path* in D if P is an induced subgraph of D. For paths P and P', by $P \circ P'$ we denote the composition of P and P', that is, the path obtained by appending P' after P. For paths $P, P_1, P_2, \ldots, P_q$ such that $P = P_1 \circ P_2 \circ \ldots \circ P_q$, we say that $P_1 \circ P_2 \circ \ldots \circ P_q$ is a *partition* of P. For a digraph D and $X \subseteq V(D)$, we say that a (u, v)-path P in D is *X-free* if none of the internal vertices of P are from X. The *X-based partition of P* in D is the partition $P = P_1 \circ \ldots \circ P_q$ such the union of the end-points of P_i, $i \in [q]$, is exactly the set $(X \cap V(P)) \cup \{u, v\}$. A *semi-$X$-based partition* of P, $P = P_1 \circ \ldots \circ P_q$, is such that the end-points of the paths P_i, $i \in [q]$, are a subset of $(X \cap V(P)) \cup \{u, v\}$. Paths $\{P_1, \ldots, P_q\}$ are *internally vertex-disjoint* if for all distinct $i, j \in [q]$, the sets of internal vertices of P_i and P_j are disjoint.

Vertex and Arc Cuts: For a digraph D and $u, v \in V(D)$, a (u, v)-*arc cut* is a set of arcs of D, say X, such that $D - X$ has no (u, v)-path. A (u, v)-*vertex cut* is a set of vertices of D, say Y, such that $D - Y$ has no (u, v)-path and $u, v \notin Y$ if $(u, v) \notin E(D)$.

2.1 Simple Results and Observations

We first give some simple results concerning the class $\mathcal{D}_\alpha$ that will be used throughout the paper.

▶ **Observation 2.1.** *Let $D \in \mathcal{D}_\alpha$. The length of the shortest cycle in D is at most $2\alpha + 1$. Also, the length of any induced path in D is at most $2\alpha + 1$.*

▶ **Lemma 2.1.** *If $D \in \mathcal{D}_\alpha$, then $|E(D)| \geq (\frac{n}{\alpha} - 1)\frac{n}{2}$.*

Proof. The proof follows from Turan's theorem [27], which states that any graph on n vertices that does not contain a clique of size $\alpha + 1$ has at most $(1 - \frac{1}{\alpha})\frac{n^2}{2}$ edges. ◀

The next lemma allows us to prove the existence of paths of any length modulo some integer p between a pair of vertices with large connectivity.

▶ **Lemma 2.2.** *Let $D \in \mathcal{D}_\alpha$ be a digraph and p be a positive integer. For $s, t \in V(D)$, if $\mathcal{P}$ is a collection of $p^2\alpha$ internally vertex-disjoint (s, t)-paths in D, then for each $i \in \{0, \ldots, p-1\}$, there exists a (s, t)-path of length $i \mod p$ in $D[V(\mathcal{P})]$.*

Proof. By the pigeonhole principle, there exist $p\alpha$ paths in $\mathcal{P}$ of the same length (without loss of generality, say 0) modulo p. Let $P_1, \ldots, P_{p\alpha}$ denote these paths. For each $j \in [p\alpha]$, let v_j be the vertex of P_j that appear after s in P_j. Since $\cup_{i \in [p\alpha]} P_i$ is a collection of internally vertex-disjoint paths, the set $X = \{v_1, \ldots, v_{p\alpha}\}$ is a set of $p\alpha$ vertices. Since $D[X] \in \mathcal{D}_\alpha$, it means that $\chi(D[X]) \geq p$ and thus, by Gallai-Roy Theorem [37, 56], there exists a path of length $p - 1$ in $D[X]$. Without loss of generality, let $P = (v_1, v_2, \ldots, v_p)$ be this path. Then for each $i \in [2..p]$, consider the path Q_i obtained as follows. Let R_1 be the (s, v_1)-subpath of P_1, that is R_1 is the arc (s, v_1). Let R_2 be (v_1, v_i)-subpath of P and let R_3 be the (v_i, t)-subpath of P_i. Then, $Q_i = R_1 \circ R_2 \circ R_3$. Clearly, length of Q_i is $i - 1 \mod p$. ◀

The following proposition relates the absence of directed cycles in a strongly connected digraph with its absence in its undirected counterpart. This will be crucially used during the NP-hardness proof of DEOCT on tournaments.

▶ **Proposition 2.3.** *If D is a strongly connected digraph with no odd directed cycles, then the underlying undirected graph of D has no odd cycles, that is, it is bipartite.*

Proof. For the sake of contradiction, suppose that D is not bipartite, that is, it contains an undirected odd cycle, say C. Let $C = (v_0, \ldots, v_{t-1})$. If for all $i \in [t-1]$, $(v_i, v_{i+1}) \in E(D)$ (here addition is modulo t), then C is also a directed odd cycle in D, which is a contradiction. Otherwise, consider any pair v_i, v_{i+1}, such that $(v_{i+1}, v_i) \in E(D)$. Since D is a strongly connected digraph, there is path, say P_i from v_i to v_{i+1} in D. If P_i is an even path then P_i together with the arc (v_{i+1}, v_i) is an odd cycle in D (which is again a contradiction). Otherwise, for all pairs v_i, v_{i+1} such that $(v_{i+1}, v_i) \in E(D)$, there is an odd (v_i, v_{i+1})-path P_i. For each pair v_i, v_{i+1} such that $(v_{i+1}, v_i) \in E(D)$, replace the arc $(v_{i+1}, v_i) \in E(D)$ with the path P_i, in C, to obtain a directed closed odd walk. Since every directed closed odd walk contains a directed odd cycle, we conclude that D contains an odd cycle, which is a contradiction. ◀

3 Finding Small k-Cut Preserving Sets

We give the precise definition of a k-cut preserving set here.

▶ **Definition 3.1** (k-Cut Preserving Set). *For digraph D, an ordered pair (u, v) of vertices of D and a positive integer k, $\{u, v\} \subseteq \mathcal{Z} \subseteq V(D)$ is a* k-cut preserving set for (u, v) in D *if the following holds. For any (u, v)-path P in D, there exists a semi-$\mathcal{Z}$-based partition $P_1 \circ \ldots \circ P_d$ of P with the following two properties. For each $i \in [d]$, P_i is an (s_i, t_i)-path in D with $s_i, t_i \in \mathcal{Z}$. Moreover, either $V(P_i) \subseteq \mathcal{Z}$ or there exists a list $\mathcal{L}_i$ of $k+1$ internally vertex-disjoint $(V(D) \setminus \mathcal{Z})$-free (s_i, t_i)-paths. A list $\mathcal{L}_i$ with the above property is called a* replacement kit for P_i *in $\mathcal{Z}$. Such a semi-$\mathcal{Z}$-based partition of P is called a* $\mathcal{Z}$-replacement witness for P.

Before moving to the computational aspects of a k-cut preserving set, we give the following lemma that can be considered as the main utility of k-cut preserving sets, and relate to the intuition we gave in the previous section.

▶ **Lemma 3.2.** *Let D be a digraph, $u, v \in V(D)$ and $\mathcal{Z}$ be a k-cut preserving set for (u, v) in D. For any set $A \subseteq E(D)$ of at most k arcs, if there exists a (u, v)-path in $D - A$, then there also exists one in $D[\mathcal{Z}] - A$.*

Proof. Consider some $A \subseteq E(D)$ such that $|A| \leq k$. Suppose there exists a (u, v)-path P in $D - A$. Since $\mathcal{Z}$ is a k-cut preserving set for the pair (u, v), there exists a semi-$\mathcal{Z}$-based partition $P = P_1 \circ \ldots \circ P_d$ such that for each $j \in [d]$, P_j is an (s_j, t_j)-path, $s_j, t_j \in \mathcal{Z}$ and, either $V(P_j) \subseteq \mathcal{Z}$, in which case P_j is a path in $D[\mathcal{Z}] - A$, or there exist $k+1$ internally vertex- disjoint (s_j, t_j)-paths in $D[\mathcal{Z}]$. In the later case, at least one of the $k+1$ paths is in $D[\mathcal{Z}] - A$ (because $|A| \leq k$). This implies the existence of a walk from u to v (and hence also a (u, v)-path) in $D[\mathcal{Z}] - A$. This concludes the proof. ◀

The main goal of this section is to prove the following lemma.

▶ **Lemma 3.3** (k-Cut Preserving Lemma). *Let D be an acyclic digraph, and $u, v \in V(D)$ be such that $N^-(u) = N^+(v) = \emptyset$. Additionally, let $D - \{u, v\} \in \mathcal{D}_\alpha$. Then there exists a k-cut preserving set for (u, v) in D of size at most $f(\alpha)$, where $f(1) = k^3 + 5k^2 + 3k$ and for $\alpha > 1$, $f(\alpha) = k^2 g(\alpha) + 2kh(\alpha)$, $g(\alpha) = (2k + (k + kf(\alpha - 1))^2) f(\alpha - 1)$ and $h(\alpha) = (k^2 + k)g(\alpha) + kf(\alpha - 1)$. Moreover, such a set can be found in time $n^{\mathcal{O}(1)}$, where $n = |V(D)|$.*

Note that $V(D)$ is always a k-cut preserving set for any pair of vertices (u, v) in D, for any k. We now define a notation, for the sake of convenience, that will be used throughout this

section. For any digraph D, $u, v \in V(D)$ and $X \subseteq V(D)$, let $ver_D(u, v; X)$ denote the union of the sets of vertices of all X-free (u, v)-paths in D. Observe that $ver_D(u, v; X) \cap X \subseteq \{u, v\}$. We begin by making an observation that forms the base line for computing small sized k-cut preserving sets using an appropriate induction.

▶ **Observation 3.1.** *Let D be a digraph, $u, v \in V(D)$, $\mathcal{Z} \subseteq V(D)$ and k be a positive integer. Let P be a (u, v)-path in D, and $P = P_1 \circ \ldots \circ P_d$ be a semi-$\mathcal{Z}$-based partition of P. If for each $i \in [d]$, there is a $\mathcal{Z}_i$-replacement witness for P_i in D_i, for some $\mathcal{Z}_i \subseteq \mathcal{Z}$ and D_i subgraph of D, then there is a $\mathcal{Z}$-replacement witness for P.*

Proof. For each $i \in [d]$, let $P_I = P_{i,1} \circ \ldots \circ P_{i,c_i}$ be a $\mathcal{Z}_i$-replacement witness for $\mathcal{Z}_i$ in D_i. Then, consider the semi-$\mathcal{Z}$-based partition $P = P_{1,1} \circ \ldots \circ P_{1,c_1} \circ P_{2,1} \circ \ldots \circ P_{2,c_2} \circ \ldots \circ P_{d,1} \circ \ldots \circ P_{d,c_d}$. Then, for each $i \in [d]$ and $j \in [c_i]$, either $V(P_{i,j}) \subseteq \mathcal{Z}_i \subseteq \mathcal{Z}$, or there exists a list $\mathcal{Z}_{i,j}$ containing $k+1$ internally vertex-disjoint $(V(D_i) \setminus \mathcal{Z}_i)$-free $(x_{i,j}, y_{i,j})$-paths in D_i such that $P_{i,j}$ is a $(x_{i,j}, y_{i,j})$-path. Since $\mathcal{Z}_i \subseteq \mathcal{Z}$ and D_i is a subgraph of D, the paths in $\mathcal{L}_{i,j}$ are $(V(D) \setminus \mathcal{Z})$-free and exist in D. ◀

Next, we give two lemmas (Lemmas 3.4 and 3.5) that basically use Observation 3.1 in a more concrete setting required to prove the k-Cut Preserving Lemma by induction on the size of the maximum independent set in the digraph.

▶ **Lemma 3.4.** *Let D be a digraph, $u, v \in V(D)$ and k be a positive integer. Let C be some (u, v)-vertex cut in D. For each $c \in C$, let $\mathcal{Z}(u, c)$ (resp. $\mathcal{Z}(c, v)$) be a k-cut preserving set for (u, c) (resp. (c, v)) in $D[ver_D(u, c; C)]$ (resp. $D[ver_D(c, v; C)]$). For each $(c, c') \in C^2$, $c \neq c'$, let $\mathcal{Z}(c, c')$ be a k-cut preserving set for (c, c') in $D[ver_D(c, c'; C)]$. Then, $\mathcal{Z} := \bigcup_{c \in C}(\mathcal{Z}(u, c) \cup \mathcal{Z}(c, v)) \cup \bigcup_{(c,c') \in C^2, c \neq c'} \mathcal{Z}(c, c')$ is a k-cut preserving set for (u, v) in D.*

Proof. First observe, from the definition of a k-cut preserving set and the construction of $\mathcal{Z}$, that $C \subseteq \mathcal{Z}$. Consider any (u, v)-path P in D. Let $P = P_1 \circ \ldots \circ P_q$ be the C-based partition of P. Since $C \subseteq \mathcal{Z}$, $P_1 \circ \ldots \circ P_q$ is a semi-$\mathcal{Z}$-based partition of P. Then P_1 is a C-free (u, c_1)-path in D for some $c_1 \in C$, P_q is a C-free (c_2, v)-path in D for some $c_2 \in C$, and for each $i \in [2, q-1]$, P_i is a C-free $(c_{j_i}, c_{j_i'})$-path in D, for some $c_{j_i}, c_{j_i'} \in C$, $j_i \neq j_i'$. Thus, P_1 is a (u, c_1)-path in $D[ver_D(u, c_1; C)]$, P_q is a (c_2, v)-path in $D[ver_D(c_2, v; C)]$, and for each $i \in [2, q-1]$, P_i is a $(c_{j_i}, c_{j_i'})$-path in $D[ver_D(c_{j_i}, c_{j_i'}; C)]$. Since $\mathcal{Z}(u, c_1), \mathcal{Z}(c_2, v), \cup_{i \in [2,q-1]} \mathcal{Z}(c_{j_i}, c_{j_i'}) \subseteq \mathcal{Z}$, we are done by Observation 3.1. ◀

▶ **Lemma 3.5.** *Let D be a digraph, $u, v \in V(D)$, and k be a positive integer. Let C be some (u, v)-vertex cut in D. For each $c \in C$, let $\mathcal{Z}(u, c)$ (resp. $\mathcal{Z}(c, v)$) be a k-cut preserving set for (u, c) (resp. (c, v)) in $D[ver_D(u, c; C)]$ (resp. $D[ver_D(c, v; C)]$). Let $X = N_D^-(v) \cap \bigcup_{c \in C} \mathcal{Z}(c, v)$. For each $(\mathtt{a}, \mathtt{b}) \in (C \cup X)^2$, $\mathtt{a} \neq \mathtt{b}$, let $\mathcal{Z}(\mathtt{a}, \mathtt{b})$ be a k-cut preserving set for $(\mathtt{a}, \mathtt{b})$ in $D[ver_D(\mathtt{a}, \mathtt{b}; C \cup N_D^-(v))]$. Then, $\mathcal{Z} := \bigcup_{c \in C}(\mathcal{Z}(u, c) \cup \mathcal{Z}(c, v)) \cup \bigcup_{(\mathtt{a},\mathtt{b}) \in (C \cup X)^2, \mathtt{a} \neq \mathtt{b}} \mathcal{Z}(\mathtt{a}, \mathtt{b})$ is a k-cut preserving set for (u, v) in D.*

Proof. First observe that $\{u, v\} \cup C \cup X \subseteq \mathcal{Z}$. Let $Y = N_D^-(v) \setminus X$. We begin by defining some special types of paths (see Figure 2).

1. A path P *is of Type* $(u, \square)$ *(resp.* $(\square, v)$*)* if it is a C-free (u, c)-path (respectively (c, v)-path) in D for some $c \in C$.
2. A path P *is of Type* $(\boxtimes, \boxtimes)$ if it is a $(C \cup N_D^-(v))$-free $(\mathtt{a}, \mathtt{b})$-path in D for some $(\mathtt{a}, \mathtt{b}) \in (C \cup X)^2$.

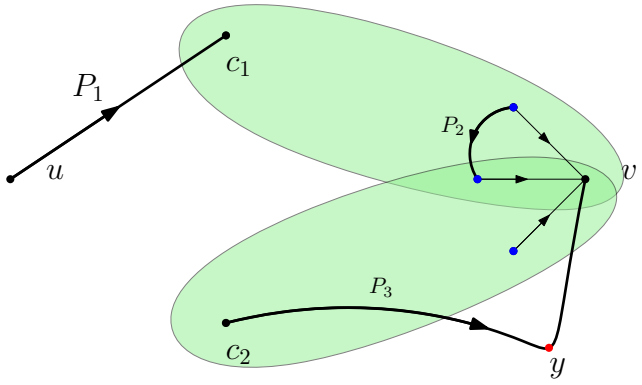

Figure 2 (c_1, c_2) is a (u, v) vertex-cut, the green parts correspond to the $\mathcal{Z}(c_i, v)$ and the blue vertices are the vertices of X. P_1 is a path of Type $(u, \square)$, P_2 is a path of Type $(\boxtimes, \boxtimes)$ and P_3 is a path of Type $(\square, \boxminus, v)$ with $y \in Y$.

3. A path P *is of Type* $(\square, \boxminus, v)$ if it is a (c, v)-path in D for some $c \in C$ and there exists $y \in V(P) \cap Y$ such that the (c, y)-subpath of P is C-free.[3]

We now begin with the proof of the lemma. Let P be some (u, v)-path. We need to show that there is a $\mathcal{Z}$-replacement witness for P. Let $P = P'_1 \circ \ldots \circ P'_q$ be the $(C \cup X)$-based partition of P. If P is not Y-free, that is, $V(P) \cap Y \neq \emptyset$, let $s' \in [q]$ be the least integer such that $V(P'_{s'}] \cap Y \neq \emptyset$. If P is Y-free, let $s' = q$. Let $s \leq s'$ be the largest integer such that P_s is an $(\mathtt{a}, \mathtt{b})$-path, where $\mathtt{a} \in C$ and $\mathtt{b} \in C \cup X \cup \{v\}$. WE first show that such a s always exists. From the definition of s', either there exists some $y \in Y$ in $V(P'_{s'})$ or $v \in V(P'_{s'})$. In the later case, since C is a (u, v)-vertex cut, there exists $c \in C$ such that c appears on P. Since $P = P'_1 \circ \ldots \circ P'_q$ is a $C \cup X$-based partition of P, there exists $s \leq s'$ such that P_s is a $(\mathtt{a}, \mathtt{b})$-path where $\mathtt{a} \in C$. In the former case again, since $y \in Y \subseteq N_D^-(v)$ and C is a (u, v)-vertex cut using previous arguments the existence of the desired s is guaranteed.

Consider the partition $P = P_1 \circ \ldots \circ P_s$, such that $P_i = P'_i$, if $i < s$ and $P_s = P'_s \circ P'_{s+1} \circ \ldots \circ P'_q$. Observe that, since $C \cup X \subseteq \mathcal{Z}$, $P = P_1 \circ \ldots \circ P_s$ is a semi-$\mathcal{Z}$-based partition of P.

▷ Claim 3.6. P_1 is a Type $(u, \square)$ path, for each $i \in [2, s-1]$, P_i is a Type $(\boxtimes, \boxtimes)$ path and, P_s is either a Type $(\square, v)$ or Type $(\square, \boxminus, v)$ path.

Proof. Recall that $P = P'_1 \circ \ldots \circ P'_q$ is the $(C \cup X)$-based partition of P. Thus, we have the following.

1. For each $i \in [q]$, P'_i is $(C \cup X)$-free path.
2. For each $i \in [2, q-1]$, P'_i is a $(\mathtt{a}, \mathtt{b})$-path, where $(\mathtt{a}, \mathtt{b}) \in (C \cup X)^2$.
3. Since C is a (u, v)-vertex cut in D and $X \subseteq N_D^-(v)$, P'_1 is a (u, c)-path for some $c \in C$.
4. From the choice of s, for each $i \in [s-1]$, $V(P'_i) \cap Y = \emptyset$. Since for $i \in [s-1]$, $P_i = P'_i$ and $X \cup Y = N_D^-(v)$, P_i is $(C \cup N_D^-(v))$-free.

Thus, from Points 2 and 4, for each $i \in [s-1]$, P_i is of Type $(\boxtimes, \boxtimes)$. Also, from Points 3 and 4, P_1 is of Type $(u, \square)$. We now show that P_s is of Type $(\square, v)$ or $(\square, \boxminus, v)$. From the

[3] Specifically, if there exists $y \in V(P) \cap Y$ with this property, then the first vertex of P that belongs to Y also has that property.

choice of s and the construction of P_s, P_s is a (c, v)-path for some $c \in C$. If P is Y-free, then P_s is of Type $(\square, v)$, otherwise, P_s is of Type $(\square, \boxminus, v)$. ◀

For each $i \in [s]$, define $\mathcal{Z}_i$ and D_i as follows.

$$\mathcal{Z}_i = \begin{cases} \mathcal{Z}(u,c) & \text{if } i = 1,\ P_1 \text{ is a } (u,c)\text{-path}, c \in C \\ \mathcal{Z}(\mathtt{a},\mathtt{b}) & \text{if } i \in [2, s-1],\ P_i \text{ is a } (\mathtt{a},\mathtt{b})\text{-path}, (\mathtt{a},\mathtt{b}) \in (C \cup X)^2 \\ \mathcal{Z}(c,v) & \text{if } i = s,\ P_s \text{ is a } (c,v)\text{-path}, c \in C \end{cases}$$

$$D_i = \begin{cases} D[ver_D(u, c; C)] & \text{if } i = 1, P_1 \text{ is a } (u,c)\text{-path}, c \in C \\ D[ver_D(\mathtt{a},\mathtt{b}; (C \cup N_D^-(v))] & \text{if } i \in [2, s-1],\ P_i \text{ is a } (\mathtt{a},\mathtt{b})\text{-path}, (\mathtt{a},\mathtt{b}) \in (C \cup X)^2 \\ D[ver_D(c, v)] & \text{if } i = s,\ P_s \text{ is a } (c,v)\text{-path}, c \in C \end{cases}$$

Recall the construction of $\mathcal{Z}$ from the lemma statement. Observe that for each $i \in [s]$, $\mathcal{Z}_i \subseteq \mathcal{Z}$. From Observation 3.1, to give a $\mathcal{Z}$-replacement witness for P, it is enough to give a $\mathcal{Z}_i$-replacement witness for each P_i, in D_i, $i \in [s]$. Thus, the following claim will finish the proof of the lemma.

▷ Claim 3.7. For each $i \in [s]$, P_i has a $\mathcal{Z}_i$-replacement witness in D_i.

Proof. We prove the claim using the following cases.

- **Case $i = 1$:** From Claim 3.6, P_1 is a C-free (u, c)-path in D for some $c \in C$. Thus, P_1 is a (u, c)-path in D_1. Since Z_1 is a k-cut preserving set for (u, c) in D_i, there exists a $\mathcal{Z}_1$-replacement witness for P_1 in D_1.
- **Case $i \in [2, s-1]$:** From Claim 3.6, when $i \in [2, s-1]$, then P_i is a $(C \cup N_D^-(v))$-free $(\mathtt{a}, \mathtt{b})$-path in D for some $(\mathtt{a}, \mathtt{b}) \in (C \cup X)^2$. Thus, P_i is an $(\mathtt{a}, \mathtt{b})$-path in D_i. Since Z_i is a k-cut preserving set for $(\mathtt{a}, \mathtt{b})$ in D_i, there exists a $\mathcal{Z}_i$-replacement witness for P_i in D_i.
- **Case $i = s$:** From Claim 3.6, P_s is of either Type $(\square, v)$ or Type $(\square, \boxminus, v)$.
  - P_s **is of Type** $(\square, v)$**:** From the definition of Type $(\square, v)$, P_s is a C-free (c, v)-path in D, for some $c \in C$. Thus, P_s is a (c, v)-path in D_s. Since $\mathcal{Z}_s$ is a k-cut preserving set for (c, v) in D_s, there exists a $\mathcal{Z}_s$-replacement witness for P_s in D_s.
  - P_s **is of Type** $(\square, \boxminus, v)$**:** From the definition of Type $(\square, \boxminus, v)$, P_s is a (c, v)-path in D, for some $c \in C$, and there exists $y \in V(P) \cap Y$ such that the (c, y)-subpath of P is C-free. Let $P_s^\dagger$ be the (c, y)-subpath of P. Recall that $Y = N_D^-(v) \setminus X$. Consider the (c, v)-path in D, denoted by $\widetilde{P_s}$, obtained by appending the arc (y, v) at the end of $P_s^\dagger$. That is, $\widetilde{P_s} = P_s^\dagger \circ (y, v)$. Since $P_s^\dagger$ is a C-free path, so is $\widetilde{P_s}$. Thus $\widetilde{P_s}$ is a (c, v)-path in D_s. Since $\mathcal{Z}_s$ is a k-cut preserving set for (c, v) in D_s, there exists a semi-$\mathbb{Z}_s$-based partition of $\widetilde{P_s}$ which is a $\mathcal{Z}_s$-replacement witness for $\widetilde{P_s}$ in D_s. Let $\widetilde{P_s} = \widetilde{P_{s,1}} \circ \ldots \circ \widetilde{P_{s,r}}$ be one such partition. Since $y \in Y = N_D^-(v) \setminus X$ and $\mathcal{Z}_s \subseteq X$, $y \notin \mathcal{Z}_s$. Thus, y is an internal vertex of $\widetilde{P_{s,r}}$. Let $\widetilde{P_{s,r}}$ be an (x, v)-path. Clearly, $x \in \mathcal{Z}_s$ because $\widetilde{P_s} = \widetilde{P_{s,1}} \circ \ldots \circ \widetilde{P_{s,r}}$ is a semi-$\mathcal{Z}_s$-based partition. Let $P_{s,r}^\dagger$ be the (x, v)-subpath of $P_{s,r}$. We claim that $P_s = \widetilde{P_{s,1}} \circ \ldots \circ \widetilde{P_{s,r-1}} \circ P_{s,r}^\dagger$ is a semi-$\mathcal{Z}_s$-based partition of P_s and is also a $\mathcal{Z}_s$-replacement witness for P_s in D_s. It is clear from the discussion above that $P_s = \widetilde{P_{s,1}} \circ \ldots \circ \widetilde{P_{s,r-1}} \circ P_{s,r}^\dagger$ is a semi $\mathcal{Z}_s$-based partition of P_s. We will now show that it is a $\mathcal{Z}_s$-replacement witness for P_s in D_s.
    Since $\widetilde{P_s} = \widetilde{P_{s,1}} \circ \ldots \circ \widetilde{P_{s,r}}$ is a $\mathcal{Z}_s$-replacement witness for $\widetilde{P_s}$, we have that for each $j \in [r]$, either $V(\widetilde{P_{s,j}}) \subseteq \mathcal{Z}_s$ or there exists a list $\mathcal{L}_j$ containing $k + 1$ vertex disjoint

paths from the start vertex of $\widetilde{P_{s,j}}$ to its end vertex. Also, since $y \notin \mathcal{Z}_s$ and y is an internal vertex of $\widetilde{P_{s,r}}$, $V(\widetilde{P_{s,r}}) \not\subseteq \mathcal{Z}_s$. Thus, there is a list $\mathcal{L}_r$ containing $k+1$ vertex disjoint (x,v)-paths (recall x and v are the start and end vertices, respectively, of $\widetilde{P_{s,r}}$). Since $P_s = \widetilde{P_{s,1}} \circ \ldots \circ \widetilde{P_{s,r-1}} \circ P^{\dagger}_{s,r}$, and $P^{\dagger}_{s,r}$ is an (x,v)-path, from the above discussion for each $j \in [r-1]$, either $V(\widetilde{P_{s,j}}) \subseteq \mathcal{Z}_s$ or there exists a list $\mathcal{L}_j$ containing $k+1$ vertex disjoint paths from the start vertex of $\widetilde{P_{s,j}}$ to its end vertex. Also, there exists a list, $\mathcal{L}_r$, containing $k+1$ vertex disjoint paths from the start vertex of $P^{\dagger}_{s,r}$ to its end vertex. This completes the proof of the claim.

◀

As argued earlier, this completes the proof of the lemma. ◀

3.1 Finding a Small k-Cut Preserving Set for a Pair with Large Flow

As explained in Section 1.2, the proof of Lemma 3.3 will distinguish whether there is a k vertex-cut for (s,t) or not. The case where there is a no k vertex-cut is the easiest one, and will be dealt with the following lemma by simply keeping $k+1$ vertex disjoint paths.

▶ **Lemma 3.8.** *Let $D \in \mathcal{D}_\alpha$ be an acyclic digraph and $u, v \in V(D)$ be such that each (u,v)-vertex cut in D has size at least $k+1$. Then, a k-cut preserving set for (u,v) in D of size at most $(2\alpha-1)(k+1)+2$ exists and is computable in $n^{\mathcal{O}(1)}$ time, where $n = |V(D)|$.*

Proof. Since every (u,v)-vertex cut in D has size at least $k+1$, from Menger's Theorem, there are at least $k+1$ vertex-disjoint (u,v)-paths in D. Let $Q'_1, \ldots, Q'_{k+1}$ be a collection of some $k+1$ of these paths. We will now obtain a collection of $Q_1, \ldots Q_{k+1}$ vertex disjoint paths where the length of each Q_i is at most $2\alpha+1$. To this end, we define each Q_i as some shortest (u,v)-path using the vertices of $V(Q'_i)$. We first claim that the length of Q_i is at most $2\alpha+1$. For the sake of contradiction, suppose not. Then, from Observation 2.1, there exist $x, y \in V(Q_i)$ such that $(x,y) \in E(D)$. Since D is acyclic, x appears before y in the path Q_i. This contradicts that Q_i is a shortest (u,v)-path in $V(Q'_i)$. Let $\mathcal{Z} = \bigcup_{i \in [k+1]} V(Q_i)$. Clearly, $\{u,v\} \subseteq \mathcal{Z}$ and $|\mathcal{Z}| \leq (2\alpha-1)(k+1)+2$. The size bound follows because the length of each Q_i is at most $2\alpha+1$, and u, v are the vertices common in each Q_i. To show that $\mathcal{Z}$ is a k-cut preserving set for (u,v) in D, consider the semi-$\mathcal{Z}$-based partition of P that is P itself. Then, $\{Q_1, \ldots, Q_{k+1}\}$ is the list for P containing $k+1$ internally vertex-disjoint $(V(D) \setminus \mathcal{Z})$-free (u,v)-paths. ◀

3.2 Finding a Small k-Cut Preserving Set of a Pair in a Tournament

As explained before, the proof of Lemma 3.3 will use induction on α. The next lemma handles the base case where $\alpha = 1$. It is somewhat more complicated compared to the arguments in Section 1.2; the reason for the complication is that we consider the digraph D such that the $D - \{u,v\} \in \mathcal{D}_\alpha$. Thus D is not "exactly" a tournament. This is required in the inductive case for the proof of Lemma 3.3.

▶ **Lemma 3.9.** *Let D be an acyclic digraph. Let $u, v \in V(D)$ be such that $N^+(u) = N^-(v) = \emptyset$ and $D - \{u,v\}$ is a tournament. Then, a k-cut preserving set for (u,v) in D of size at most $k^3 + 5k^2 + 3k$ exists and is computable in polynomial time.*

Proof. If all (u,v)-vertex cuts in D have size at least $k+1$, then the correctness follows from Lemma 3.8. Thus, for the rest of the proof assume that there is a (u,v)-vertex-cut in D of size at most k. Let $C = \{c_1, \ldots, c_\ell\}$ be a minimal (u,v)-vertex cut in D of size $\ell \leq k$.

▷ Claim 3.10. $C \subseteq N_D^+(u) \cup N_D^-(v)$.

Proof. Suppose not. Then, there exists $c_i \in C$ such that $c_i \notin N^+(u) \cup N^-(v)$. Since C is a minimal (u,v)-vertex cut in D, there exists a path, say P, from u to v in $D-(C \setminus \{c_i\})$. Let u' be the first vertex on P after u and v' be the last vertex of P before v. Since $D-\{u,v\}$ is an acyclic tournament, $(u',v') \in E(D)$. Since $u', v' \notin C$, we get a (u,v)-path in $D-C$, contradicting that C is a (u,v)-vertex cut in D. ◀

Let $I = \{i \in [\ell] \mid c_i \in N_D^-(v)\}$ and $J = \{j \in [\ell] \mid c_j \in N_D^+(u)\}$. For all $i \in I$, let $U_i = ver_D(u, c_i; C)$ and $D_i = D[U_i]$. For all $j \in J$, let $V_j = ver_D(c_j, v; C)$ and $D_j = D[V_j]$. For all $(i,j) \in [\ell]^2$, $i \neq j$, let $Q_{i,j} = ver_D(c_i, c_j; \emptyset)$ and $D_{i,j} = D[Q_{i,j}]$.

For each $i \in I$ (resp. $j \in J$, resp. $(i,j) \in [\ell]^2$, $i \neq j$), we will compute a k-cut preserving set $\mathcal{Z}_i$ (resp. $\mathcal{Z}_j$, resp. $\mathcal{Z}_{i,j}$) of (u, c_i) (resp. (c_j, v), resp. (c_i, c_j)) in D_i (resp. D_j, resp. $D_{i,j}$) of size at most $2k+3$ (resp. $2k+3$, resp. $k+3$). The procedure to do so is as follows.

- **Computing $\mathcal{Z}_i$, $i \in I$:** First observe that U_i is a candidate for $\mathcal{Z}_i$. Thus, if $|U_i| \leq 2(k+1)$, set $\mathcal{Z}_i = U_i$. Otherwise, we have that $|U_i| \geq 2k+3$. Since $D-\{u,v\}$ is an acyclic tournament, let π be the unique topological ordering of $D-\{u,v\}$. We divide this case further into two cases.
  - **Case 1: $|N^+(u) \cap U_i| \leq k$:** Let $\widetilde{U_i}$ be the last $k+1$ vertices of U_i in π. Observe that $\widetilde{U_i} \subseteq N^-(c_i) \cap U_i$. Define $\mathcal{Z}_i = (N^+(u) \cap U_i) \cup \widetilde{U_i} \cup \{u, c_i\}$. Clearly, $|\mathcal{Z}_i| \leq 2k+3$. To prove that $\mathcal{Z}_i$ is a k-cut preserving set for (u, c_i) in D_i, consider some (u,c_i)-path P in D_i, such that $V(P) \not\subseteq \mathcal{Z}_i$. We will show a $\mathcal{Z}_i$-replacement witness for P in D_i. Consider the semi-$\mathcal{Z}_i$-based partition of P, $P = P_1 \uplus P_2$, where P_1 is the arc $(u,x) \in E(P)$, for some $x \in N^+(u) \cap U_i$ and P_2 is the (x,c_i)-subpath of P. Clearly, $V(P_1) \subseteq \mathcal{Z}_i$. We claim that there are $k+1$ vertex-disjoint (x,c_i)-paths in $\mathcal{Z}_i$. To see this, consider the following argument. Since $V(P) \not\subseteq \mathcal{Z}_i$, there exists a vertex $y \in V(P)$ such that $y \notin \mathcal{Z}_i$. Then, $y \in V(P_2)$. Since $y \notin \mathcal{Z}_i$, it in particular holds that $y \notin \widetilde{U_i}$. Thus, all the vertices of $\widetilde{U_i}$ appear after y in π. Since there is a (x,y)-path in D_i, x appears before y in π. Thus, x appears before all the vertices of $\widetilde{U_i}$ in π. Thus, because $D-\{u,v\}$ is a tournament, $\widetilde{U_i} \subseteq N^+(x) \cap U_i$. Since $\widetilde{U_i} \subseteq N^-(c_i) \cap U_i$, there are $|\widetilde{U_i}|$ many vertex disjoint (x,c_i)-paths in $\mathcal{Z}$. This completes the proof.
  - **Case 2: $|N^+(u) \cap U_i| > k$:** First observe that all the vertices of $N^+(u) \cap U_i$ appear before c_i in π. Since π is a topological ordering of $D-\{u,v\}$, there are $|N^+(u) \cap U_i| > k$ vertex-disjoint (u,c_i)-paths in $\mathcal{Z}_i$. Thus, each (u,c_i)-vertex-cut in D_i has size at least $k+1$. In this case, let $\mathcal{Z}_i$ be the k-cut preserving set for (u,c_i) in D_i obtained from Lemma 3.8. Observe that $|\mathcal{Z}_i| \leq k+3$.
- **Computing $\mathcal{Z}_j$, $j \in J$:** $\mathcal{Z}_j$ can be computed using arguments symmetric to the previous case.
- **Computing $\mathcal{Z}_{i,j}$, $(i,j) \in [\ell]^2$, $i \neq j$:** First observe that all the vertices of $Q_{i,j} \setminus \{c_i, c_j\}$ appear after c_i and before c_j in π. Thus, there are $|Q_{i,j} \setminus \{c_i, c_j\}|$ many vertex-disjoint (c_i, c_j)-paths in $D_{i,j}$. If $|Q_{i,j}| \leq k-2$, then set $\mathcal{Z}_i = Q_{i,j}$, otherwise let $\mathcal{Z}_i$ be the k-cut preserving set for (c_i, c_j) in $D_{i,j}$ obtained from Lemma 3.8. In either case, $|\mathcal{Z}_i| \leq k+3$.

Let $\mathcal{Z} := \bigcup_{i \in I} \mathcal{Z}_i \cup \bigcup_{j \in J} \mathcal{Z}_j \cup \bigcup_{(i,j) \in [\ell]^2, i \neq j} \mathcal{Z}_{i,j}$. Observe that $C \subseteq \mathcal{Z}$. First note that $|\mathcal{Z}| \leq |I|(2k+3) + |J|(2k+3) + \ell^2(k+3) \leq k^3 + 5k^2 + 3k^2$ (the last inequality holds because $|I| + |J| = \ell$ and $\ell \leq k$). We will now show that $\mathcal{Z}$ is a k-cut preserving set for (u,v) in D. To see this, consider some(u,v)-path P, in D. Since C is a (u,v)-vertex-cut in D there exists a vertex of C on P. Let c_i be the first vertex of C on P and c_j be the last vertex of C on P

(c_i could be the same as c_j). Let P_1 be the (u, c_i)-subpath of P, P_2 be the (c_i, c_j)-subpath of P and P_3 be the (c_j, v)-subpath of P (if c_i is the same as c_j, then P_2 is empty). Thus, $P = P_1 \circ P_2 \circ P_3$ is a semi-$\mathcal{Z}$-based partition of P (as $C \subseteq \mathcal{Z}$). Since $\mathcal{Z}_i$ is a k-cut preserving set for (u, c_i) in D_i, $\mathcal{Z}_i$ is a k-cut preserving set for (c_j, v) in D_j and $\mathcal{Z}_{i,j}$ is a k-cut preserving set for (c_i, c_j) in $D_{i,j}$, and $\mathcal{Z}_i, \mathcal{Z}_j, \mathcal{Z}_{i,j} \subseteq \mathcal{Z}$, from Observation 3.1, $\mathcal{Z}$ is a k-cut preserving set for (u, v) in D. ◀

3.3 Finding a small k-cut preserving set for a pair in a $D \in \mathcal{D}_\alpha$

We are now ready to prove Lemma 3.3.

▶ **Lemma 3.3** (k-Cut Preserving Lemma). *Let D be an acyclic digraph, and $u, v \in V(D)$ be such that $N^-(u) = N^+(v) = \emptyset$. Additionally, let $D - \{u, v\} \in \mathcal{D}_\alpha$. Then there exists a k-cut preserving set for (u, v) in D of size at most $f(\alpha)$, where $f(1) = k^3 + 5k^2 + 3k$ and for $\alpha > 1$, $f(\alpha) = k^2 g(\alpha) + 2kh(\alpha)$, $g(\alpha) = (2k + (k + kf(\alpha - 1))^2) f(\alpha - 1)$ and $h(\alpha) = (k^2 + k)g(\alpha) + kf(\alpha - 1)$. Moreover, such a set can be found in time $n^{\mathcal{O}(1)}$, where $n = |V(D)|$.*

Proof. We prove this lemma using induction on α. When $\alpha = 1$, the proof follows from Lemma 3.9.

▷ Claim 3.11. Let $x, y \in V(D) \setminus \{x, y\}$. Then, a k-cut preserving set for (x, y) of size $g(\alpha)$ in any digraph D' that is a subgraph of D where $u, v \notin V(D')$, can be found in polynomial time.

Proof. Let W be a minimum (x, y)-vertex-cut in D'. If $|W| > k$, then the claim follows from Lemma 3.8. Thus, we are now in the case where $|W| \leq k$. For each $w \in W$, let $\mathcal{Z}(x, w)$ (resp. $\mathcal{Z}(w, y)$) be a k-cut preserving set for (x, w) (resp. (w, y)) in $D'[ver_{D'}(x, w; W)]$ (resp. $D'[ver_{D'}(w, y; W)]$). Let $B = N^-_{D'}(y) \cap \bigcup_{w \in W} \mathcal{Z}(w, y)$. For each $(\mathtt{a}, \mathtt{b}) \in (W \cup B)^2$, let $\mathcal{Z}(\mathtt{a}, \mathtt{b})$ be a k-cut preserving set for $(\mathtt{a}, \mathtt{b})$ in $D'[ver'_D(\mathtt{a}, \mathtt{b}; W \cup N^-(y))]$. Then, from Lemma 3.5, $\mathcal{Z}(x, y) := \bigcup_{w \in W} (\mathcal{Z}(x, w) \cup \mathcal{Z}(w, y)) \cup \bigcup_{(\mathtt{a},\mathtt{b}) \in (W \cup B)^2} \mathcal{Z}(\mathtt{a}, \mathtt{b})$ is a k-cut preserving set for (x, y) in D'.

We will now show that for any $w \in W$ and $(\mathtt{a}, \mathtt{b}) \in (W \cup B)^2$, each digraph among $D'[ver_{D'}(x, w; W)]$, $D'[ver_{D'}(w, y; W)]$ and $D'[ver_{D'}(\mathtt{a}, \mathtt{b}; W \cup N^-_{D'}(y))]$ has independence number strictly smaller than α. Then, from induction hypothesis and the expression for $\mathcal{Z}(x, y)$ written above, we will conclude that a k-cut preserving set for (x, y) in D' of size $g(\alpha)$ can be found in polynomial time. To see that the independence number of $D'[ver_{D'}(x, w; W)]$ is strictly less than α, observe that y is not adjacent to any vertex in $ver_{D'}(x, w; W)$, as W is an (x, y)-vertex cut in D'. Thus, any independent set of $D'[ver_{D'}(x, w; W)]$ together with y is an independent set of D' and hence of D. Since $y \notin \{u, v\}$, $u, v \notin V(D')$ and the independence number of $D - \{u, v\}$ is α, we have that the independence number of $D'[ver_{D'}(x, w; W)]$ is strictly smaller than α. A similar argument holds for $D'[ver_{D'}(w, y; W)]$ as in this case x is not adjacent to any vertex of $ver_{D'}(w, y; W)$. For $D'[ver_{D'}(\mathtt{a}, \mathtt{b}; W \cup N^-_{D'}(y))]$, since $ver_{D'}(\mathtt{a}, \mathtt{b}; W \cup N^-_{D'}(y)) \cap N^-_{D'}(y) = \emptyset$, $u, v \notin V(D')$ and $N^+_{D'}(y) = \emptyset$, any independent set of $D'[ver_{D'}(\mathtt{a}, \mathtt{b}; W \cup N^-_{D'}(y))]$ together with y is an independent set in $D - \{x, y\}$. Since $D - \{x, y\}$ has independence number α, $D'[ver_{D'}(\mathtt{a}, \mathtt{b}; W \cup N^-_{D'}(y))]$ has independence number strictly smaller than α. ◀

Let C be a minimum (u, v)-vertex-cut in D. If $|C| > k$, then the lemma follows from Lemma 3.8. Thus, for the remainder of the proof we assume that $|C| \leq k$. For each $c \in C$, let $U_c = ver_D(u, c; C)$, $V_c = ver_D(c, v; C)$, $\mathcal{Z}(u, c)$ be a (u, c) k-cut preserving set in

$D[U_c]$, and $\mathcal{Z}(c,v)$ be a (c,v) k-cut preserving set in $D[V_c]$. For each $(c,c') \in C^2$, $c \neq c'$, let $Q_{c,c'} = ver_D(c,c';C)$, and $\mathcal{Z}(c,c')$ be a k-cut preserving set in $D[Q_{c,c'}]$. Then from Lemma 3.4, $\mathcal{Z} := \bigcup_{c\in C} \mathcal{Z}(u,c) \cup \mathcal{Z}(c,v) \cup \bigcup_{(c,c')\in C^2, c\neq c'} \mathcal{Z}(c,c')$ is a k-cut preserving set for (u,v) in D. Since $C \cap \{u,v\} = \emptyset$, from Claim 3.11, for each $(c,c') \in C^2$, $c \neq c'$, $\mathcal{Z}(c,c')$ of size $g(\alpha)$ can be computed in polynomial time. In the remainder of the proof, we will show how to compute $\mathcal{Z}(u,c)$ and $\mathcal{Z}(c,v)$, for any $c \in C$, of the desired size. We will only give the proof of construction of $\mathcal{Z}(u,c)$ as the proof for $\mathcal{Z}(c,v)$ is symmetrical.

▷ Claim 3.12. For any $c \in C$, $\mathcal{Z}(u,c)$ of size $h(\alpha)$ can be computed in polynomial time.

Proof. For ease of notation, let $\widehat{D} = D[U_c]$. Let A be a minimum (u,c)-vertex-cut in $\widehat{D}$. First note that $A \cap \{u,v\} = \emptyset$. If $|A| > k$, then the claim follows from Lemma 3.8. Thus, for the remainder of the proof, assume that $|A| \leq k$.

For each $a \in A$, let $\widehat{U_a} = ver_{\widehat{D}}(u,a;A)$, $\widehat{V_a} = ver_{\widehat{D}}(a,c;A)$, $\widehat{\mathcal{Z}}(u,a)$ be a (u,a) k-cut preserving set in $\widehat{D}[\widehat{U_a}]$ and $\widehat{\mathcal{Z}}(a,c)$ be a (a,c) k-cut preserving set in $\widehat{D}[\widehat{V_a}]$. For each $(a,a') \in A^2$, $a \neq a'$, let $R_{a,a'} = ver_{\widehat{D}}(a,a';A)$ and $\widehat{\mathcal{Z}}(a,a')$ be a k-cut preserving set in $\widehat{D}[R_{a,a'}]$. Then from Lemma 3.4, $\mathcal{Z}(u,c) := \bigcup_{a\in A}(\widehat{\mathcal{Z}}(u,a) \cup \widehat{\mathcal{Z}}(a,c)) \cup \bigcup_{(a,a')\in A^2, a\neq a'} \widehat{\mathcal{Z}}(a,a')$ is a k-cut preserving set for (u,c) in D. Since $A \cap \{u,v\} = \emptyset$ and $c \in \{u,v\}$, from Claim 3.11, for each $a \in A$, $(a,a') \in A^2$, $a \neq a'$, $\widehat{\mathcal{Z}}(a,c)$ and $\widehat{\mathcal{Z}}(a,a')$ of size $g(\alpha)$ can be computed in polynomial time. Moreover, the independence number of $\widehat{D}[\widehat{U_a}] - \{u,a\}$ is strictly smaller than α because $c(\neq v)$ is not adjacent to any vertex in $\widehat{U_a}$, besides possibly u and a. Thus, for each $a \in A$, a set $\widehat{\mathcal{Z}}(u,a)$ of size $f(\alpha-1)$ can be computed in polynomial time by the induction hypothesis. This finishes the proof of the claim. ◀

Thus, from the previous arguments and Claim 3.12, we have that $\mathcal{Z}$ is a k-cut preserving set for (u,v) in D of size at most $k^2 g(\alpha) + 2kh(\alpha)$. ◀

A rough computation gives that, for any k, $g(\alpha) \leq 6k^2 f(\alpha-1)$ and $h(\alpha) \leq 8k^4 f(\alpha-1)$. This imply that $f(\alpha) \leq 22k^5 f(\alpha-1)^3$. By noting that $f(1) \leq 22k^5$, we can show the following observation.

▶ **Observation 3.13.** *For any α and k, there exists a k-cut preserving set of size smaller than $f(k,\alpha) = (22k^5)^{4^\alpha}$.*

3.4 k-Cut Preserving Sets for a Set of Vertices

Below we also define a notion of k-cut preserving sets for a set of vertices. Such a notion will come handy in our applications. Given a digraph D and $X \subseteq V(D)$, for each $(u,v) \in X^2$, we define the digraph $D^X_{(u,v)}$ as follows (note that u could be equal to v). Let $R = V(D) - X$. Then, $D^X_{(u,v)}$ is the supergraph of $D[R]$ obtained by adding two new vertices u^+ and v^- together with the following set of additional arcs: $\{(u^+,x) : x \in R, (u,x) \in E(D)\} \cup \{(x,v^-) : x \in R, (x,v) \in E(D)\}$.

▶ **Definition 3.14** (k-Cut Preserving Set for a Set of Vertices)**.** *For any digraph D, a positive integer k and $X \subseteq V(D)$, we say that $X \subseteq \mathcal{Z} \subseteq V(D)$ is a* k-cut preserving set for X*, if for all $(u,v) \in X^2$, $\mathcal{Z}$ is a k-cut preserving set for (u,v) in $D^X_{(u,v)}$.*

▶ **Lemma 3.15.** *For any digraph $D \in \mathcal{D}_\alpha$, a positive integer k, and $S \subseteq V(D)$ such that $D - S$ is a acyclic, a k-cut preserving set for S of size at most $|S|^2 f(k,\alpha)$ can be found in polynomial time, where $f(k,\alpha) \leq (22k^5)^{4^\alpha}$.*

Proof. For each pair $(u,v) \in S^2$ (u and v could be equal), let $\mathcal{Z}_{(u,v)}$ be the a k-cut preserving set for (u^+, v^-) in $D^S_{(u,v)}$ obtained from Lemma 3.3. From the definition of k-cut preserving set for S, $\mathcal{Z} = \bigcup_{(u,v)\in S^2} \mathcal{Z}_{(u,v)}$ is a k-cut preserving set for S. From Observation 3.13, for any $(u,v) \in S^2$, $|\mathcal{Z}_{(u,v)}| \leq f(k,\alpha)$. Thus, we conclude the correctness of the lemma. ◀

4 Applications of the k-Cut Preserving Lemma

4.1 Fault-Tolerant (S,S)-Reachability

In this section, we prove Lemma 1.1. Recall that (D, S, ℓ, k) is an instance of FTR(S,S) where $D \in \mathcal{D}_\alpha$, $S \subseteq V(D)$ and ℓ, k are positive integers such that each strongly connected component of $D - S$ has size at most ℓ. The goal is to compute a subgraph H of D of size $k^{2^{\mathcal{O}(\alpha)}}$ such that, for any $A \subseteq E(D)$ of size at most k, for any $s,t \in S$, if $D - A$ has an (s,t)-path, then so does $H - A$. It is not difficult to see from Lemma 3.2 that if $\mathcal{Z}$ is a k-cut preserving set for S in D, then $H = D[\mathcal{Z}]$ is a solution for (D, S, ℓ, k) (for any ℓ). When $\ell = 1$, $D - S$ is acyclic and hence a k-cut preserving set for S can be computed using Lemma 3.15. When $\ell > 1$, in order to use Lemma 3.15 we modify the digraph D to turn $D - S$ acyclic. We now describe the operation, which we call dagify, that is used to turn $D - S$ acyclic. Informally, for each strongly connected component SC of D we turn it into an independent set while preserving the paths in D that use the vertices of SC. This is achieved by creating a new vertex for every ordered pair of vertices (say, (u,v)) in SC. Such a vertex represents the existence of a (u,v)-path in the strongly connected component SC. In fact, in the path in the modified graph, each new vertex corresponding to some pair (u,v) can be replaced by some (u,v)-path from the strongly connected component SC to yield a path in the original graph. Then, arcs between two vertices in this newly constructed vertex set are put in such a way that the concatenation of the paths corresponding to these new vertices gives a path in D. This idea is formalized below.

▶ **Definition 4.1** (dagify(D,R)). *Let D be a digraph, $R \subseteq V(D)$ and $S = V(D) \setminus R$. Let $SC_1, \ldots, SC_d$ be the strongly connected components of $D[R]$. For $a \in [d]$, let $V(SC_a) = \{v^a_1, \ldots, v^a_{n_a}\}$, where $n_a = |V(SC_a)|$. Then, $D^\dagger_R :=$ dagify(D,R) is the digraph defined as:*

Vertex set of $D^\dagger_R$: *For each $a \in [d]$, let $SC^\dagger_a = \{\mathbf{v}^a_{ij} \mid (v^a_i, v^a_j) \in \{SC_a\}^2, i,j \in [n_a]\}$. Let $R^\dagger = \cup_{a\in[d]} SC^\dagger_a$ and $V(D^\dagger_R) = R^\dagger \cup S$.*

Arc set of $D^\dagger_R$: *It contains all the arcs of D with both end-points in S. For each $a \in [d]$, $SC^\dagger{}_a$ is an independent set in $D^\dagger_R$. For any $a \in [d]$, $s \in S$ and $i,j \in [n_a]$, $(s, \mathbf{v}^a_{ij}) \in E(D^\dagger_R)$ if and only if $(s, v^a_i) \in E(D)$. Similarly, $(\mathbf{v}^a_{ij}, s) \in E(D^\dagger_R)$ if and only if $(v^a_j, s) \in E(D)$. We put the arcs between $SC^\dagger_a$ and $SC^\dagger_b$, for distinct $a,b \in [d]$ as follows. For any $i,j \in [n_a]$ and $i',j' \in [n_b]$, $(\mathbf{v}^a_{ij}, \mathbf{v}^b_{i'j'}) \in E(D^\dagger_R)$ if and only if $(v^a_j, v^b_{i'}) \in E(D)$.*

For a set of vertices of $X^\dagger \subseteq D^\dagger_R$, *full-comp*$(X^\dagger)$ denotes the set of vertices of $V(D)$ such that, for each $\mathbf{v}^a_{i,j} \in X^\dagger$, all the vertices of SC_a belong to full-comp$(X^\dagger)$. Also all the vertices of S that belong to $X^\dagger$, belong to *full-comp*$(X^\dagger)$. Observe that $|$*full-comp*$(X^\dagger)| \leq \ell^2 \cdot |X^\dagger|$, where ℓ is the upper bound on the size of each SC_a. Note from the construction above that, for any $s,t \in S$ and an (s,t)-path $P^\dagger$ in $D^\dagger_R$, there exists an (s,t)-path P in D such that $V(P) \subseteq$ *full-comp*$(P^\dagger)$. The following observations state a few properties of the digraph $D^\dagger_R$ that would be useful when we want to find a k-cut preserving set for $D^\dagger_R$ using Lemma 3.15.

▶ **Observation 4.1.** *$D^\dagger_R[R^\dagger]$ is acyclic.*

Proof. Recall, from the construction of $D^\dagger_R$, that $R^\dagger = \bigcup_{a\in[d]} SC^\dagger_a$ and each $SC^\dagger_a$ is an independent set in $D^\dagger_R$. Without loss of generality, let $SC_1, \ldots, SC_d$ be the strongly connected components of $D[R]$ ordered as in their topological ordering. Then, there is no arc from a vertex of SC_b to a vertex of SC_a, for any $b > a$, in D. Thus, from the construction of $D^\dagger_R$, there is no arc from any $\mathbf{v}^b_{ij}$ to any $\mathbf{v}^a_{i'j'}$ $(b > a)$. This shows that $D^\dagger_R[R^\dagger]$ is acyclic. ◀

▶ **Observation 4.2.** *If $D \in \mathcal{D}_\alpha$ and every strongly connected component of $D[R]$ has size at most ℓ, then $D^\dagger_R \in \mathcal{D}_{\ell^2\alpha}$.*

Proof. Recall that $R^\dagger = \bigcup_{a\in[d]} SC^\dagger{}_a$ and $D^\dagger_R[SC^\dagger_a]$ has no arc. From the construction of $D^\dagger_R$, for each $a \in [d]$, $|SC^\dagger_a| \leq \ell^2$. Finally, since $D \in \mathcal{D}_\alpha$, from the construction of $D^\dagger_R$, the size of any maximum independent set in $D^\dagger_R$ is at most $\max_{a\in[d]} |SC^\dagger_a| \cdot \alpha \leq \ell^2\alpha$. ◀

We define some terminology that would come handy later. For any $A \subseteq E(D)$, we say that a vertex $v \in V(D)$ *is affected by* A if there exists some arc of A that is incident on v. The *set affected by A in $D^\dagger_R$* is the set of vertices of $D^\dagger_R$ containing the union of the vertices in $SC^\dagger_a$, for each $a \in [d]$ such that a vertex in SC_a is affected by A in D.

▶ **Observation 4.3.** *Let D be a digraph, $R \subseteq V(D)$ and $S = V(D) \setminus R$. Let $A \subseteq E(D)$ of size at most k. Let $\mathcal{A}^\dagger$ be the set affected by A in $D^\dagger_R$. Recall the construction of $D^\dagger_R$ from Definition 4.1. For some $\mathbf{v}^a_{ij}, \mathbf{v}^b_{i'j'} \in R^\dagger$, let $P^\dagger$ be an $\mathcal{A}^\dagger$-free $(\mathbf{v}^a_{ij}, \mathbf{v}^b_{i'j'})$-path in $D^\dagger_R$. Then there exists a $(v^a_i, v^b_{j'})$-path P in D such that: $V(P) \subseteq full\text{-}comp(P^\dagger)$ and, P does not use any arc of A.*

Proof. Recall the construction of $\mathsf{dagify}(D, R)$. Consider any path P obtained from $P^\dagger$ by replacing all the vertices of $R^\dagger$ as follows. If for any $c \in [d]$, $i^*, j^* \in [n_c]$, $\mathbf{v}^c_{i^*j^*} \in V(P^\dagger)$, then replace $\mathbf{v}^c_{i^*j^*}$ in $P^\dagger$ by any $(v^c_{i^*}, v^c_{j^*})$-path in the strongly connected component SC_c. Clearly, the path P obtained is a $(v^a_i, v^b_{j'})$-path in D and $V(P) \subseteq full\text{-}comp(P^\dagger)$. Also from the definition of $\mathcal{A}^\dagger$ and the fact that $P^\dagger$ is $\mathcal{A}^\dagger$-free, we get that P cannot use an arc of A. ◀

From the construction in Definition 4.1, for any $s, t \in S$, for an (s,t)-path P in D, we can associate a unique (s,t)-path $P^\dagger$ in $D^\dagger_R$. This is elaborated below. Consider the digraph $D^\dagger_R$ obtained by $\mathsf{dagify}(D, R)$. $(v^a_i, v^a_j) \in SC^2_a$ for some component SC_a of $D[R]$. Let $s, t \in S$. Let P be an S-free (s,t)-path in D. For any such path P, we define the notion of a *reduced path* of P in $D^\dagger_R$ as follows. Consider the unique partition $P = P_s \circ P_{i_1} \circ \ldots \circ P_{i_q} \circ P_t$ such that P_s is an arc (s,u) where $u \in V(SC_{i_1})$, P_t is an arc (v,t) where $v \in V(SC_{i_q})$ and for each $j \in [q]$, $V(P_{i_j}) \subseteq V(SC_{i_j})$, where $i_1, \ldots, i_j \in [d]$ and $i_1 < \ldots < i_q$. For each $j \in [q]$, let P_{i_j} be a $(v^{i_j}_{p_j}, v^{i_j}_{r_j})$-path. Consider the vertex $\mathbf{v}^{i_j}_{p_j,r_j}$ in $V_{i_j} \subseteq R^\dagger \subseteq V(D^\dagger_R)$. From the construction of $D^\dagger_R$, we get the (s,t)-path $P^\dagger = s \circ \mathbf{v}^{i_1}_{p_1,r_1} \circ \mathbf{v}^{i_2}_{p_2,r_2} \circ \ldots \circ \mathbf{v}^{i_q}_{p_q,r_q} \circ t$ in $D^\dagger_R$. This (s,t)-path $P^\dagger$ in $D^\dagger_R$ is called the *reduced path of P in $D^\dagger_R$*.

Proof of Lemma 1.1. Recall (D, S, ℓ, k) is an instance of $\mathrm{FTR}(S, S)$. Let $R = V(D) \setminus S$. Let $D^\dagger_R$ be obtained by $\mathsf{dagify}(D, R)$. From Observations 4.1 and 4.2, Lemma 3.15 can be used to compute a $(2k\ell^2 + 1)$-cut preserving set for S in $D^\dagger_R$. Let $\mathcal{Z}^\dagger$ be such a set. Let $\mathcal{Z} = full\text{-}comp(\mathcal{Z}^\dagger)$. We claim that $H = D[\mathcal{Z}]$ is a solution to the instance (D, S, ℓ, k). (First note that the size bound on H follows from Lemma 3.15 and the fact that each strongly connected component of R has size at most ℓ.)

Towards this let $A \subseteq E(D)$ of size at most k, $s, t \in S$ and P be an (s,t)-path in $D - A$. We need to show that there is some (s,t)-path in $H - A$ too. Let $P = P_1 \circ \ldots \circ P_q$ be the S-based partition of P such that each P_i is an (s_i, t_i)-path. Then it suffices to show that for each fixed $i \in [q]$, there is some (s_i, t_i) path in $H - A$ (these paths would yield a closed

walk from s to t in $H - A$ and hence an (s,t)-path in $H - A$). In the remaining part of the proof, we focus on proving this. Note that each P_i is S-free. Fix any $i \in [q]$. For the ease of notation, let us call the path P_i as P, vertices s_i, t_i as s, t respectively.

Let $P^\dagger$ be the reduced path corresponding of P in $D^\dagger_R$. Since $\mathcal{Z}^\dagger$ is a $(2k\ell^2+1)$-cut preserving set for $P^\dagger$ in $D^\dagger_R$, consider a $\mathcal{Z}^\dagger$-witnessing replacement $P^\dagger = P^\dagger_1 \circ \ldots \circ P^\dagger_r$. Recall the notation from the construction in Definition 4.1.

For an arbitrary $c \in [r]$, let $P^\dagger_c$ be a $(\mathbf{v}^a_{ij}, \mathbf{v}^b_{i',j'})$-path (or $(s, \mathbf{v}^a_{ij})$-path or $(\mathbf{v}^a_{ij}, s)$-path). Observe that, since $P^\dagger$ is the reduced path of P, to finish the proof of the lemma, it is enough to show a (v^a_i, v^b_j)-path (or (s, v^a_i)-path or (v^a_i, s)-path) exists in $H - A$. Without loss of generality, let $P^\dagger_c$ be a $(\mathbf{v}^a_{ij}, \mathbf{v}^b_{i',j'})$-path, the other cases hold due to similar arguments.

As $P^\dagger = P^\dagger_1 \circ \ldots \circ P^\dagger_d$ is a $\mathcal{Z}^\dagger$-witnessing replacement, one of the following cases arises.

1. $V(P^\dagger_c) \subseteq \mathcal{Z}^\dagger$. Since $P^\dagger$ is the reduced path of P, consider the (v^a_i, v^b_j)-subpath, say P'_c, of P. Then, $V(P'_c) \subseteq \mathit{full\text{-}comp}(P^\dagger_c) \subseteq \mathcal{Z}$ (because $V(P^\dagger_c) \subseteq \mathcal{Z}^\dagger$). Also since P does not have an arc in A, so does P'_c. Thus, by the construction of H, P'_c is a path in $H - A$.
2. There is a list $\mathcal{L}_i$ of $2k\ell^2+1$ internally vertex-disjoint $(\mathbf{v}^a_{ij}, \mathbf{v}^b_{i'j'})$-paths in $D^\dagger_R[\mathcal{Z}^\dagger]$. Let $\mathcal{A}^\dagger$ be the set of affected vertices of A in $D^\dagger_R$. Clearly, $|\mathcal{A}^\dagger| \leq 2k\ell^2$. Then there exists a path in $\mathcal{L}_i$ that is $\mathcal{A}^\dagger$-free. Then from Observation 4.3, there exists a (v^a_i, v^b_j)-path, say P'_c, such that $V(P'_C) \subseteq \mathit{full\text{-}comp}(P^\dagger_c) \subseteq \mathcal{Z}$ and, that does not use an arc of A. From the construction of H, P'_c is a path in $H - A$.

This finishes the proof of the lemma. ◀

4.2 Kernel for DFAS on $\mathcal{D}_\alpha$

In this section, we give a polynomial kernel for DFAS on $\mathcal{D}_\alpha$, that is, we prove Theorem 1.3.

▶ **Theorem 1.3.** DFAS *on* $\mathcal{D}_\alpha$ *admits a kernel of size* $k^{\mathcal{O}(4^\alpha)}$.

We achieve this in two steps. In the first step, we find a set of vertices S of size $\mathcal{O}(\alpha k)$ whose removal results in an acyclic digraph. We then show that it is enough to keep the vertices of a k-cut-preserving set for S to get a kernel.

▶ **Lemma 4.2.** *Let* (D,k) *be an instance of* DFAS *and let* $D \in \mathcal{D}_\alpha$. *In polynomial time, one can either correctly conclude that* (D,k) *is a* NO *instance of* DFAS, *or output a set* $S \subseteq V(D)$ *such that* $|S| \leq (2\alpha+1)k$ *and* $D - S$ *is acyclic.*

Proof. Since $D \in \mathcal{D}_\alpha$, if there exists a cycle in D, then from Observation 2.1 there exists a cycle of length at most $2\alpha+1$. Thus, one can greedily find vertex-disjoint cycles each of length at most $2\alpha+1$. To see this, notice that after the removal of any vertex set, the resulting digraph remains in $\mathcal{D}_\alpha$ and hence our previous argument reapplies. If one finds more than k such cycles, then any *dfas* of S has size at least $k+1$, in which case report that (D,k) is a NO instance. Otherwise, one finds a collection $\mathcal{C}$ of at most k cycles of length at most $2\alpha+1$ each such that every cycle of D intersects in some vertex of one of the cycles in $\mathcal{C}$. In this case, output S as the union of the vertex sets of the cycles in $\mathcal{C}$. ◀

▶ **Lemma 4.3.** *Let* (D,k) *be an instance of* DFAS *where* $D \in \mathcal{D}_\alpha$. *Let* S *be a set computed by Lemma 4.2 on input* (D,k). *Let* $\mathcal{Z}$ *be a* k*-cut preserving set for* S *in* D *computed by Lemma 3.15. Then,* (D,k) *is a* YES *instance of* DFAS *if and only if* $(D[\mathcal{Z}],k)$ *is a* YES *instance.*

Proof. Since $D[\mathcal{Z}]$ is a subgraph of D, the forward direction is trivial. For the backward direction, let A be a *dfas* of $D[\mathcal{Z}]$ of size at most k. We will prove that A is also a *dfas* of D. For the sake of contradiction, suppose that A is not a *dfas* of D, that is, there is some cycle C in $D - A$. Since $D - S$ is acyclic, C must contain some vertex from S. Let $v_0, \ldots, v_\ell$ be the vertices in $V(C) \cap S$, appearing in this order along C (the choice of which vertex is denoted v_0 is arbitrary).

For any pair (v_i, v_{i+1}), the (v_i, v_{i+1})-subpath of C is S-free. Recall the construction of the digraph $D^S_{(v_i,v_{i+1})}$ before Definition 3.14. From the definition of $\mathcal{Z}$ (by Definition 3.14), $\mathcal{Z}$ contains a k-cut preserving set for (v_i, v_{i+1}) in $D^S_{v_i,v_{i+1}}$, for each $i \in [0, \ell]$. From Lemma 3.2, there exists a path P'_i from v_i to v_{i+1} in $D[\mathcal{Z}] - A$. Thus, we conclude that for each $i \in [\ell]_0$, there exists a (v_i, v_{i+1})-path (where addition is modulo ℓ) in $D[\mathcal{Z}] - A$. Since C is a cycle, these paths give a closed walk (and hence also a cycle) in $D[\mathcal{Z}] - A$. ◀

Proof of Theorem 1.3. Its correctness follows from Lemmas 3.15 and 4.3 by noting that the size of the set $\mathcal{Z}$ obtained is smaller than $((2\alpha + 1)k)^2 f(k, \alpha)$. ◀

5 Introducing Parity Preserving Set

One of the mains tools employed for the proof of Lemma 1.1 and Theorem 1.3 was Lemma 3.2, which says that if $\mathcal{Z}$ is a k-cut preserving set for (u, v) and A is a set of at most k arcs, then the existence of a (u, v)-path in $D - A$ implies the existence of one in $D[\mathcal{Z}] - A$. To prove Lemma 1.2 and Theorem 1.7, we need to take into account not only the existence of a path, but also its length modulo a certain integer p. For this reason, we introduce a notion of parity preserving sets. Vaguely speaking, in the a parity preserving set, like the cut preserving set, we need to keep witnesses for paths. In the case of cut preserving sets keeping a list of $k + 1$ paths was enough to ensure the existence of one of them after the removal of at most k arcs, as one path from each of the lists, together would yield a walk which was sufficient for the cut-preserving purposes. In this case, since a walk doesn't necessarily yield the required parity path, we need to have witness lists of large enough size that ensure the existence of enough paths such that a disjoint witness of required parity can be found for each piece to yield a witnessing path for the given path. Also, the size of the list of paths required as a witness then also becomes a function of the size of the original path. This is formalized in the definition below.

▶ **Definition 5.1** ((k, p, q, t)-parity preserving set for a collection of pairs)**.** *For any digraph D, positive integers k, p, q, t, and a collection of pairs of vertices $\mathcal{V} \subseteq V(D)^2$, a set $\mathcal{Z}$ is called a* (k, p, q, t)-parity preserving set for $\mathcal{V}$ *if for any $A \subseteq E(D)$ of size at most k, if there exists $\mathcal{V}' \subseteq \mathcal{V}$, $|\mathcal{V}'| \leq q$, $\mathcal{V}' = \{(u_i, v_i) : i \in [q'], q' \leq q\}$ such that,*

1. *for each $(u_i, v_i) \in \mathcal{V}'$, there is a (u_i, v_i)-path P_i of length at most t in $D - A$,*
2. *for each $i, j \in [q']$, $i \neq j$, the paths P_i, P_j described above are internally vertex-disjoint,*
3. *for each $i \in [q']$, the internal vertices of P_i are disjoint from the set of vertices in the pairs of $\mathcal{V}$,*

then for each $i \in [q']$, there exists a (u_i, v_i)-path P^*_i in $D[\mathcal{Z}] - A$ such that,

1. *length of P_i modulo p is equal to length of P^*_i modulo p,*
2. *for each $i, j \in [q']$, $i \neq j$, the paths P_i and P^*_j are internally vertex-disjoint.*

We now show how to a find (k,p,q,t)-parity preserving sets using cut preserving sets. For this, first recall the definition of $dagify(D,R)$ (Definiton 4.1). Recall that, the strongly connected components of R are denoted by SC_i and their corresponding vertices in $R^\dagger$ are denoted by $SC_i^\dagger$.

▶ **Lemma 5.2.** *Let k,p,q,t,ℓ be positive integers. Let $D \in \mathcal{D}_\alpha$ be a digraph, $S \subseteq V(D)$ and $R = V(D) - S$. Suppose that every strongly connected component of $D[R]$ contains at most ℓ vertices. Let $\beta = qt(2\ell^2\alpha+1)p^2\alpha\ell^2 + (qt+2k)\ell^2 + p^2\alpha$. Let $\mathcal{Z}^\dagger$ be a β-cut preserving set for S in $dagify(D,R) = D_R^\dagger$, and let $\mathcal{Z} = S \cup \{SC_i \mid$ there exists $v \in SC_i^\dagger$ such that $v \in \mathcal{Z}^\dagger\}$. Then, $\mathcal{Z}$ is a (k,p,q,t)-parity preserving set for S^2 in D. Moreover, such a set $\mathcal{Z}$ of size at most $|S|^2 f(\beta,\alpha\ell^2)\ell^2$ can be computed, where f is as defined in Lemma 3.15.*

Proof. From Observation 4.2, $D_R^\dagger \in \mathcal{D}_{\ell^2\alpha}$, which together with Lemma 3.15, implies that β-cut preserving set $\mathcal{Z}^\dagger$ for S in $D^{\dagger_R}$ of size at most $|S|^2 f(\beta,\alpha\ell^2)$ can be computed. Thus, $|\mathcal{Z}| \leq |S|^2 f(\beta,\alpha\ell^2)\ell^2$. It remains to show that $\mathcal{Z}$ is a (k,p,q,t)-parity preserving set for S^2 in D. For simplicity of notation, let $D^\dagger$ denote $D_R^\dagger$.

Recall the notation above Lemma 3.15. Notice that $\mathcal{Z}^\dagger$ is the union of β-cut preserving sets, specifically, a cut $C_{u,v}$ for (u^+,v^-) in $D^\dagger{}^S_{u,v}$ for all $u,v \in S$. For convenience, we will consider that any vertex $s \in S$ of $D_R^\dagger$ corresponds to the pair of vertices (s,s) of D. Because of this, any vertex v of $D_R^\dagger$ corresponds to a pair of vertices of D. In some arguments ahead, for a vertex (a_1,b_1) of $D_R^\dagger$, we will want to consider a (a_1,b_1)-path in the component of $D[R]$ containing a_1 and b_1. When $a_1 = b_1 = s \in S$, it will be always safe to take the single vertex s as this path.

Let $A \subseteq E(D)$ be of size at most k. Let $\{(u_1,v_1),\ldots,(u_c,v_c)\}$ be some pairs of vertices of S, where $c \leq q$ and let $P_1,\ldots,P_c$ be a set of paths such that,

1. for each $i \in [c]$, P_i is a (u_i,v_i)-path of length at most t in $D - A$,
2. for each $i,j \in [c]$, $i \neq j$, the paths P_i, P_j are internally vertex-disjoint,
3. for every $i \in [c]$, the internal vertices of P_i are disjoint from S.

We want to show the existence of a set of paths $P_1^*,\ldots,P_c^*$ such that every P_i^* is a (u_i,v_i)-path in $D[\mathcal{Z}] - A$, the length of P_i modulo p is the same as the length of P_i^* modulo p and, all the P_i^*'s are internally vertex-disjoint.

For each $i \in [c]$, let $P_i^\dagger$ be the reduced path of P_i in $D^\dagger$. Note that $P_i^\dagger$ is a (u_i,v_i)-path in $D^\dagger$. Since, $\mathcal{Z}^\dagger$ is a β-cut preserving set for (u_i,v_i) in $D^\dagger$, there exists a semi-$\mathcal{Z}^\dagger$-based partition of $P_i^\dagger$, $P_i^\dagger = P_{i,1}^\dagger \circ \cdots \circ P_{i,a_i}^\dagger$, such that for each $j \in [a_i]$, either $V(P_{i,j}^\dagger) \subseteq \mathcal{Z}^\dagger$ or there exists a list $\mathcal{L}_{i,j}^\dagger$ of size β containing paths between the same endpoints as $P_{i,j}^\dagger$.

▶ Note 5.3. For each $i \in [c]$, since the length of P_i is at most t, $a_i \leq t$.

▷ Claim 5.4. Without loss of generality, for each $i \in [c], j \in [a_i]$, all the paths in $\mathcal{L}_{i,j}^\dagger$ have length at most $2\ell^2\alpha + 1$.

Proof. Recall from Observation 4.2 that $D^\dagger \in \mathcal{D}_{\ell^2\alpha}$. Thus, from Observation 2.1, we can safely assume that all the paths in $\mathcal{L}_{i,j}^\dagger$ have length at most $2\ell^2\alpha+1$. ◀

For each path $P_{i,j}^\dagger$ in $D^\dagger$, we will now associate a path $P_{i,j}$ in D with it. The path $P_{i,j}$ is basically the subpath of P_i with the endpoints corresponding to the endpoints of $P_{i,j}^\dagger$, that is, $P_{i,j}$ is obtained from $P_{i,j}^\dagger$ by replacing each vertex $\mathbf{v}_{x,y}^z$ in $P_{i,j}^\dagger$ by some (v_x,v_y)-path in SC_z. Consider the collection of paths $\{P_{i,j} \mid i \in [c], j \in [a_i]\}$. Observe that our goal reduces to finding another collection $\{P_{i,j}^* \mid i \in [c], j \in [a_i]\}$ such that,

1. the vertex sets of the internal vertices of $P^*_{i,j}$ and $P^*_{i',j'}$ are disjoint, for each $i, i' \in [c]$, $j \in [a_i]$, $j' \in [a_{i'}]$, $i \neq i'$, $j \neq j'$,
2. for each $i \in [c], j \in [a_i]$, the length of the path $P_{i,j}$ modulo p is the same as the length of the path $P^*_{i,j}$ modulo p.

It is enough to prove the above because the collection $P^*_i = P^*_{i,1} \circ \cdots \circ P^*_{i,a_i}$, for each $i \in [c]$ is the desired collection.

If $V(P^\dagger_{i,j}) \subseteq \mathcal{Z}^\dagger$, then let $P^*_{i,j} = P_{i,j}$. Note that, in this case, $V(P_{i,j}) = V(P^*_{i,j}) \subseteq \mathcal{Z}$. Thus, $P^*_{i,j}$ is a path in $D[\mathcal{Z}] - A$. In the other case, when there exists a list $\mathcal{L}^\dagger_{i,j}$ of paths of size β for $P^\dagger_{i,j}$, we clean this list as follows.

First, since $c \leq q$, all the P_i's are of length at most t, each strongly connected component in $D[R]$ contains at most ℓ vertices and $|A| \leq k$, we need to remove only $(qt + 2k)\ell^2$ paths from each list $\mathcal{L}^\dagger_{i,j}$ so that the internal vertices of the remaining paths correspond to strongly connected components of $D[R]$ not containing any of the vertices in any of the P_i's or any vertex adjacent to any edge in A. Since the number of lists $\mathcal{L}^\dagger_{i,j}$ is bounded by ct (refer Note 5.3), and each path is smaller than $2\ell^2\alpha + 1$, we can remove at most $qt(2\ell^2\alpha + 1)p^2\alpha\ell^2$ further paths from each list so that the remaining lists, which we denote as $\widehat{\mathcal{L}}^\dagger_{i,j}$'s, consist of $p^2\alpha$ internally vertex-disjoint paths of $D\dagger^S_{u_i,v_i}$, such that if SC_r is a strongly connected component of $D[R]$ and a pair of vertices $(x_1, x_2) \in V(SC_r)$ is used in a path of $\widehat{\mathcal{L}^\dagger}_{i,j}$, then no other path of any other $\mathcal{L}_{x,q}$ can use a pair of vertices of $V(SC_r)$ as an internal vertex.

Let $Q^a_{i,j} = x_1 \circ x_2 \circ \cdots \circ x_f$ be a path of $\mathcal{L}_{i,j}$, with all the x_i's being vertices of C_{u_i,v_i}. We will construct a path $T^a_{i,j}$ of $D[\mathcal{Z}] - A$ as follows: x_1 is a vertex of $P^\dagger_i$ which corresponds in P_i to a path in some strongly connected component SC_r from some vertex a_1 to a vertex b_1. Let K_1 be this path. Likewise, x_f is a vertex of $P^\dagger_i$ which corresponds in P_i to a path in some strongly connected component $SC_{r'}$ from some vertex a_f to a vertex b_f. Let K_f be this path. Each of the other x_i's corresponds to a pair of vertices (a_i, b_i) in some strongly connected component SC_{r_i} of $D[R]$, let K_i be any (a_i, b_i)-path in this component. By definition of the digraph $D^\dagger_R$, it is clear that the concatenation of the K_i using the additional arcs from b_r to a_{r+1} for $r \in [f-1]$ is a (a_1, b_f)-path in $D[\mathcal{Z}]$. Moreover, because of the cleaning part of our argument, we know that none of the x_r for $r \in [2, f-1]$ belong to strongly connected component adjacent to arcs of S, which means that all the K_r for $r \in [2..f-1]$ are in $D[\mathcal{Z}] - A$. Since K_1 and K_f are subpaths of H_i, and don't contain any vertex not in C_{u_i,v_i}, they also belong to $D[\mathcal{Z}] - A$.

Let $Q^a_{i,j}$ be a path of some $\widehat{\mathcal{L}}^\dagger_{i,j}$, $Q^{a'}_{i',j'}$ a path of some $\mathcal{L}_{i'.j'}$ (possibly $i = i'$ and $j = j'$) and $T^a_{i,j}$ and $T^{a'}_{i',j'}$ the paths of $D[\mathcal{Z}] - A$ associated. Because of the cleaning part of our procedure, the internal vertices of $Q^a_{i.j}$ and $Q^{a'}_{i',j'}$ belong to different strongly connected components of $D[R]$. This implies that the paths $T^a_{i,j}$ and $T^{a'}_{i',j'}$ are internally vertex-disjoint. For a similar reason, they are also internally vertex disjoint from all the vertices of all the other P_j. It means that, for any fixed i, j such that $P^\dagger_{i,j}$ does not belong to C_{u_i,v_i}, the only thing we have left to argue is that there exists, among the paths $T^a_{i,j}$, a path of the same length modulo p as $P^*_{i,j}$. This is done by Lemma 2.2, as there is at least $p^2\alpha$ of those paths. This ends the proof. ◀

6 Applications of Parity Preserving Sets

In this section we show how to utilize Lemma 5.2 to prove Lemma 1.2 and Theorem 1.7.

6.1 Parity Reachability Fault Tolerance

In order to prove Lemma 1.2, we need a way to bound the size of the paths that we consider. This is the purpose of the next two lemmas.

▶ **Lemma 6.1.** *Let $D \in \mathcal{D}_\alpha$, p be some positive integer and u and v be two vertices such that any strongly connected component in $D - \{u, v\}$ has size at most ℓ. If there exists a path P from u to v of length q such that $q \mod p = r$ for some $r \in [p-1]$ and $q \geq \alpha p\ell + 2$, then there exists a path P' from u to v of length q' such that $q' \mod p = (r+1)$ and $q - \alpha p\ell \leq q' < q$. Moreover, $V(P') \subset V(P)$.*

Proof. Suppose $P = x_0, \ldots, x_{q-1}$ and consider the set of vertices $S = \cup_{i \in [0,\alpha]} x_{1+ip\ell}$. Because $D \in \mathcal{D}_\alpha$ and S is a set of size $\alpha + 1$, there is an arc between two vertices $x_{1+ip\ell}$ and $x_{1+jp\ell}$ for some $0 \leq i < j \leq \alpha$. This arc has to be oriented from $x_{1+ip\ell}$ to $x_{1+jp\ell}$ or it would create with the subpath of P from $x_{1+ip\ell}$ to $x_{1+jp\ell}$ a cycle of length greater than ℓ in $D - \{u, v\}$. Thus the path P' obtained from P by replacing the subpath of P from $x_{1+ip\ell}$ to $x_{1+jp\ell}$ by the arc $(x_{1+ip\ell}, x_{1+jp\ell})$ satisfies all the properties required. ◀

If P is a path of length at least $\alpha p^2\ell$, it means we can apply Lemma 6.1 p times to get a path of the same parity. This gives the following lemma.

▶ **Lemma 6.2.** *Let D be a digraph in $\mathcal{D}_\alpha$, p some positive integer and u and v two vertices such that any strongly connected component in $D - \{u, v\}$ has size at most ℓ. If there exists a path P from u to v of length q such that $q \mod p = r$ for some $r \in [p-1]$, then there exists a path P' from u to v of length q' such that $q' \mod p = r$ and $q' \leq \alpha p^2\ell + 2$. Moreover, $V(P') \subset V(P)$.*

We are now ready to prove the parity version of our fault-tolerant result.

▶ **Lemma 1.2.** *Given a digraph $D \in \mathcal{D}_\alpha$, positive integers k, ℓ, p, a non-negative integer r, and $S \subseteq V(D)$ such that every strongly connected component of $D - S$ has at most ℓ vertices, the* Fault-Tolerance (S,S)-Parity Reachability *problem admits a solution H on $(|S|\alpha\ell pk)^{\mathcal{O}(4^{\alpha\ell^2})}$ vertices. Furthermore, such a solution H can be found in polynomial time.*

Proof. Let $\mathcal{Z}'$ be a $(k, p, |S|, \alpha p^2\ell + 2)$-parity-preserving set obtained by applying Lemma 5.2 to D and S. A rough computation would show that the β defined in Lemma 5.2 is then smaller than $20|S|\alpha^3\ell^5p^5 + 2k\ell^2$, which gives that:

$$|\mathcal{Z}'| \leq |S|^2\ell^2(22(20|S|\alpha^3\ell^5p^5 + 2k\ell^2)^5)^{4^{\alpha\ell^2}}.$$

Let us now show that $\mathcal{Z}'$ is a solution to the Fault-Tolerance (S,S)-Parity Reachability problem. Let A be a set at most k arcs, s and t two vertices of S, and P a path from s to t in $D - A$. Let $s_1, \ldots, s_e$ denote the vertices in the intersection of P with S, in the order as they appear on P, and for every $i \in [e-1]$, P_i denote the subpath of P from s_i to s_{i+1}. As P is a path, $e \leq |S|$. By applying Lemma 6.2, we can assume that all the P_i's are smaller than $\alpha p^2\ell + 2$. Thus, by definition of a $(k, p, |S|, \alpha p^2\ell + 2)$-parity-preserving set, for every $i \in [e]$ there exists a path P'_i in $D[\mathcal{Z}'] - A$ from s_i to s_{i+1} of the same length modulo p as P_i and such that all the P'_i are internally vertex-disjoint. Taking the union of the P'_i gives the desired (s,t)-path in $D[\mathcal{Z}'] - A$. ◀

6.2 Kernel for MOD(p)-DCT in $\mathcal{D}_\alpha$

In this section, we present a proof of Theorem 1.7. The proof follows the same structure as the proof of Theorem 1.3. First we need to find an approximate solution of polynomial size. For this we need the following result, due to Chen et al. [21]

▶ **Theorem 6.3.** *Let $l \geq 2$ be an integer. If a strongly connected digraph D contains no directed cycle of length $1 \mod p$, then $\chi(D) \leq p$.*

Remember that, when $D \in \mathcal{D}_\alpha$, $\chi(D) \leq p$ implies that $|D| \leq \alpha p$.

▶ **Lemma 6.4.** *Let $D \in \mathcal{D}_\alpha$ and p a positive integer. Then either D does not contain a cycle of length $1 \mod p$, or such a cycle on fewer than $p(\alpha+1)^2$ vertices exists.*

Proof. Suppose C is the smallest cycle of length $1 \mod p$, and $|C| = q \geq (\alpha+1)^2 p$. Let $C = x_0, \ldots, x_{q-1}$ denote the vertices of C. Consider the set of vertices $A = \cup_{i \in [0,\alpha]} \{x_{i\alpha p}\}$. Because A contains more than α vertices, there is an arc between two vertices of A, say from $x_{i\alpha p}$ to $x_{j\alpha p}$. However, since $q \geq (\alpha+1)^2 p$, the subpath of C from $x_{j\alpha p}$ to $x_{i\alpha p}$ contains more than αp vertices. Let C' denote the vertices on this path. $D[C']$ is a strongly connected graph on more than αp vertices. Theorem 6.3 implies the existence of a cycle of length $1 \mod p$ in $D[C']$ which is a contradiction as C' is smaller than C. ◀

With the previous lemma, one can easily adapt the proof of Lemma 4.2 to show the following:

▶ **Lemma 6.5.** *Let (D,k) be an instance of* MOD(p)-DCT *and let $D \in \mathcal{D}_\alpha$. In polynomial time, one can either correctly conclude that (D,k) is a* NO *instance of* MOD(p)-DCT*, or output a set $S \subseteq V(D)$ such that $|S| \leq (\alpha+1)^2 pk$ and $D - S$ does not have any cycle of length $1 \mod p$.*

We are now ready to prove the existence of a kernel for MOD(p)-DCT:

▶ **Theorem 1.7.** MOD(p)-DCT *on $\mathcal{D}_\alpha$ admits a kernel of size $(p\alpha k)^{\mathcal{O}(4^{\alpha^3 p^2})}$.*

Proof. Let (D,k) be an instance MOD(p)-DCT. By applying Lemma 6.5, we can either find $k+1$ vertex disjoint cycles of length $1 \mod p$, and conclude that (D,k) is a NO instance, or find a set S of size at most $k(\alpha+1)^2 p$ vertices such that $D - S$ doesn't contain any cycle of length $1 \mod p$. Let $R = V(D) - S$ and note that Theorem 6.3 implies that the strongly connected component of $D[R]$ have at most αp vertices. Let $\mathcal{Z}$ be a $(k, p, p(\alpha+1+k)^2, p(\alpha+1+k)^2)$-parity preserving set for S obtained from applying Lemma 5.2 to D and S. Note that, the β defined in Lemma 5.2 is then smaller than $10p^8\alpha^{10}k^4$ and thus $|\mathcal{Z}| \leq (k(\alpha+1)^2 p)^2 (22(10p^8\alpha^{10}6k^4)^5)^{4^{\alpha^3 p^2}}$. We claim that (D,k) is a YES instance of MOD(p)-DCT if and only if $(D[\mathcal{Z}], k)$ is a YES instance of MOD(p)-DCT.

For the ease of notation, let $D' = D[\mathcal{Z}]$. Since D' is a subgraph of D, the forward direction is trivial. For the backward direction, let A be a set of at most k arcs such that of $D' - A$ has no cycle of length $1 \mod p$. We will now prove that $D' - A$ also has no cycle of length $1 \mod p$. For the sake of contradiction, suppose there is a cycle of length $1 \mod p$ in $D - A$ and let C be the smallest such cycle. Since $D \in \mathcal{D}_\alpha$ and $|A| \leq k$, $D - A \in \mathcal{D}_{\alpha+k}$. Then by Lemma 6.4, the length of C is at most $p(\alpha+1+k)^2$.

Since $D[R]$ has no cycle of length $1 \mod p$, $C \cap S \neq \emptyset$. Let $v_1, \ldots, v_q$ be the vertices of $C \cap S$ in the order as they appear on C. Note that $q \leq p(\alpha+1+k)^2$. Then, for each $i \in [q]$, there is a subpath $P_{i,i+1}$ (count $q+1$ as 1) from v_i to v_{i+1}, such that $P_{i,i+1}$ is S-free. Note that all the length of each of these paths is at most the length of the cycle C which is at most $p(\alpha+1+k)^2$.

▷ Claim 6.6. For each $i \in [q]$, there exists a path $P'_{i,i+1}$ from v_i to v_{i+1} in $D' - A$ such that:

1. the length modulo p of $P'_{i,i+1}$ is the same as the one of $P_{i,i+1}$, and
2. for any $i, j \in [q]$, $i \neq j$, the set of internal vertices of $P'_{i,i+1}$ and $P'_{j,j+1}$ are disjoint.

Proof. Since $\mathcal{Z}$ is a $(k, p, p(\alpha+1+k)^2, p(\alpha+1+k)^2)$-parity-preserving set for S^2 in D and $\bigcup_{i\in[q]}(v_i, v_{i+1}) \subseteq S^2$, the claim follows from the definition of $(k, p, p(\alpha+1+k)^2, p(\alpha+1+k)^2)$-parity-preserving set for S^2. ◀

Consider the cycle C' formed by taking the union of all the paths $P'_{i,i+1}$, for all $i \in [q]$. From Claim 6.6, C' exists in $D' - A$ and has the same length modulo p as C. This contradicts the definition of A and ends the proof. ◀

6.3 Difference Between Paths and Cycles

As stated in the introduction, there is a fundamental difference between Lemma 1.2 and Theorem 1.7 as we are only able to obtain a kernel for cycle of length modulo 1 mod p. The reason behind this difference is that Theorem 6.3 only works for cycles of length 1 mod p (they even provide counterexamples in [21] when the length is not 1), while Lemma 6.2 works for any $r \in [0, p-1]$.

7 NP-hardness of Directed Edge Odd Cycle Transversal on Tournaments

In this section, we will show that DEOCT is NP-hard on tournaments by showing that it is equivalent to DFAS on tournaments. Given a digraph D, observe that any *dfas* of D is also a *deoct*. But the converse may not always be true. But what we can prove in the converse case is that if $|S|$ is an *deoct* of D, then there exists a vertex set $|S'|$ such that $|S'| \leq |S|$ and S' is a *dfas* of D. Lemma 7.2 proves is. The following lemma from [55] will be used in the proof of Lemma 7.2.

▶ **Lemma 7.1** ([55] Lemma 11). *Let D be a tournament on n vertices and m arcs. Then D has a dfas of size at most $\frac{m}{2} + \frac{1}{2}\lceil\frac{n-1}{2}\rceil$.*

▶ **Lemma 7.2.** *Let D be a tournament. If S is a deoct of D, then there exists a dfas of D of size at most $|S|$.*

Proof. Consider the digraph $D - S$. Let $C_1, \ldots, C_t$ be the set of strongly connected components of $D - S$. For ease of notation, let us denote $D[C_i]$ by D_i. Let $S_i = S \cap E(C_i)$. For each $i \in [t]$, we will now construct a set X_i such that X_i is a *dfas* of D_i such that $|X_i| \leq |S_i|$. Since D_i is a strongly connected digraph and has no odd directed cycles, from Proposition 2.3, D_i is bipartite. Let (A_i, B_i) be a bipartition of D_i. Let $|A_i| = n_a$ and $|B_i| = n_b$.

▷ Claim 7.3. $|S_i| \geq \binom{n_a}{2} + \binom{n_b}{2}$.

Proof. Since (A_i, B_i) is a bipartition of D_i, $D_i[A_i]$ and $D_i[B_i]$ are independent. Since D_i is an induced subgraph of $D - S$ and D is a tournament, all the arcs whose both end-points are either in A_i or in B_i, are in S. Thus, the claim follows. ◀

Let X_i^a and X_i^b be the *dfas* of $D[A_i]$ and $D[B_i]$ obtained from Lemma 7.1. Then $|X_i^a| \leq \frac{1}{2}\binom{n_a}{2} - \frac{1}{2}\lceil\frac{n_a-1}{2}\rceil$ and $|X_i^b| \leq \frac{1}{2}\binom{n_b}{2} - \frac{1}{2}\lceil\frac{n_b-1}{2}\rceil$. Consider $E(A_i, B_i)$ and $E(B_i, A_i)$. If $|E(A_i, B_i)| \leq |E(B_i, A_i)|$, then let $X_i^{ab} = E(A_i, B_i)$. Otherwise, let $X_i^{ab} = E(B_i, A_i)$. Observe that $|X_i^{ab}| \leq \frac{n_a n_b}{2}$. Let $X_i = X_i^a \cup X_i^b \cup X_i^{ab}$.

▷ Claim 7.4. X_i is a *dfas* of D_i.

Proof. Suppose not. Then there is a cycle, say Q, in $D_i - X_i$. Recall (A_i, B_i) is a bipartition of D_i. Since X_i^a is a *dfas* of $D[A_i]$ and $X_i^a \subseteq X_i$, Q is not entirely contained in A_i. Similarly, Q is not entirely contained in B_i. Thus, if such a cycle Q exists, it has to intersect both A_i and B_i. This implies there exists two distinct arcs of Q, say e_1 and e_2, such that $e_1 \in E(A_i, B_i)$ and $e_2 \in E(B_i, A_i)$. But this is not possible, because $X_i^{ab} \subseteq X$. ◀

▷ Claim 7.5. $|X_i| \leq |S_i|$.

Proof. From the construction of X_i, we have $|X_i| = |X_i^a| + |X_i^b| + |X_i^{ab}$. Thus,

$$\begin{aligned}
|X_i| &\leq \frac{1}{2}\binom{n_a}{2} - \frac{1}{2}\lceil\frac{n_a-1}{2}\rceil + \frac{1}{2}\binom{n_b}{2} - \frac{1}{2}\lceil\frac{n_b-1}{2}\rceil + \frac{n_a n_b}{2} \\
&= \frac{1}{2}\binom{n_a}{2} - \frac{1}{2}\lceil\frac{n_a-1}{2}\rceil + \frac{1}{2}\binom{n_b}{2} - \frac{1}{2}\lceil\frac{n_b-1}{2}\rceil + \frac{n_a n_b}{2} \\
&= \frac{n_a^2}{4} - \frac{n_a}{4} - \frac{1}{2}\lceil\frac{n_a-1}{2}\rceil + \frac{n_b^2}{4} - \frac{n_b}{4} - \frac{1}{2}\lceil\frac{n_b-1}{2}\rceil + \frac{n_a n_b}{2} \\
&\leq \frac{n_a^2}{4} - \frac{n_a}{4} + \frac{1}{2}(\frac{n_a}{2} - \frac{1}{2}) - \frac{n_b^2}{4} + \frac{n_b}{4} - \frac{1}{2}(\frac{n_b}{2} - \frac{1}{2}) + \frac{n_a n_b}{2}
\end{aligned}$$

Observe that $\frac{n_a n_b}{2} \leq \frac{n_a^2}{4} + \frac{n_b^2}{4}$. Thus, we have the following.

$$|X_i| \leq \frac{n_a^2}{2} + \frac{n_b^2}{2} - \frac{n_a}{2} + \frac{n_b}{2} + \frac{1}{8} = \binom{n_a}{2} + \binom{n_b}{2} + \frac{1}{8}$$

Since, $\binom{n_a}{2} + \binom{n_b}{2}$ is an integer and the size of the set X_i is an integer, we have that $|X_i| \leq \binom{n_a}{2} + \binom{n_b}{2} \leq |S|$. The last inequality follows from Claim 7.3.

◀

Let $S' = S \setminus \cup_{i\in[t]} S_i$. Observe that $S = S' \uplus S_1 \uplus \ldots \uplus S_t$. Let $X = \cup_{i\in[t]} X_i \cup S'$.

▷ Claim 7.6. $|X| \leq |S|$.

Proof. Since $X = \cup_{i\in[t]} X_i \cup S'$, $|X| = \cup_{i\in[t]} |X_i| + |S'|$. Thus, from Claim 7.5, $|X| \leq \cup_{i\in[t]} |X_i| + |S'| \leq |S|$. ◀

▷ Claim 7.7. X is a *dfas* of D.

Proof. For the sake of contradiction, suppose there is a cycle, say Q, in $D - X$. Recall that $C_1, \ldots, C_t$ are the strongly connected components of $D - S$. Also, S' is the set of those arcs of S whose one endpoint belong to C_i and the other in C_j, for some $i, j \in [t]$, $i \neq j$. Since $S' \subseteq X$, the vertex set of Q cannot intersect both C_i and C_j for some $i, j \in [t]$, $i \neq j$. Thus, the vertex set of$Q$ is fully contained in some C_i. Since $X_i \subseteq X$ and X_i is a *dfas* of D_i (from Claim 7.4), there is no cycle in $D_i - X$. This proves the claim. ◀

Claim 7.6 and 7.7 prove the lemma. ◀

▶ **Lemma 7.8.** *Let D be a tournament. For any integer k, D has a dfas of size at most k if and only if D has a deoct of size at most k.*

Proof. Clearly, any *dfas* of D is also a *deoct* of D. The other direction follows from Lemma 7.2. ◀

Since DFAS on tournaments is NP-hard [58], from Lemma 7.8 it follows that, DEOCT on tournaments in NP-hard. This proves Theorem 1.5.

8 Sub-exponential FPT Algorithms

A k-*cut* of a digraph D is a partition of the vertex set of D into two parts, $V(D) = L \uplus R$, such that $|E(R, L)| \leq k$. Misra et al. [48] proved the following bound on the number of k-cuts in any digraph $D \in \mathcal{D}_\alpha$.

▶ **Lemma 8.1** ([48], Lemma 4). *If $D \in \mathcal{D}_\alpha$, then for any positive integer k, the number of k-cuts in D is most $2^{c\sqrt{k}\log k} \cdot (n+1)^{2\alpha\lceil\sqrt{k}\rceil} \cdot \log n$, where c is a fixed absolute constant.*

Further, the k-cuts in any digraph can be enumerated by polynomial delay.

▶ **Lemma 8.2** ([35], Lemma 7). *k-cuts of any digraph D can be enumerated with polynomial-time delay.*

8.1 Improved Sub-exponential FPT Algorithm for DFAS on $\mathcal{D}_\alpha$.

A sub-exponential FPT algorithm for DFAS was presented in [48, Theorem 1] with running time $2^{\alpha^2\sqrt{k}\log(\alpha k)}n^{\mathcal{O}(\alpha)}$. This algorithm is obtained by a dynamic programming on the number of k-cuts in an input instance (D, k). The above running time directly follows from the number of k-cuts in a digraph $D \in D_\alpha$ of bounded out-degeneracy [48, Lemma 7, Lemma 20]. We can obtain a faster algorithm by first applying Theorem 1.3 to the input instance (D, k) to obtain a kernel (D', k'), and then applying [48, Theorem 1] to (D', k'). This procedure gives a running time of $2^{\mathcal{O}(f(\alpha)\sqrt{k}\log k)} \cdot n^{\mathcal{O}(1)}$ for some function f, of α), thereby proving Theorem 1.4.

8.2 Sub-exponential FPT Algorithm for Directed Edge Odd Cycle Transversal on $\mathcal{D}_\alpha$

In this section, we prove Theorem 1.8.

▶ **Theorem 1.8.** DEOCT *on $\mathcal{D}_\alpha$ admits an algorithm with running time $2^{\mathcal{O}(f(\alpha)\sqrt{k}\log k)} \cdot n^{\mathcal{O}(1)}$, where $f(\alpha)$ is a function of α and n is the number of vertices in D.*

Our approch is similar to [48], but requires some additional work to handle the even cycles that remain after removing a solution.

For a digraph D, let us define a γ-*vertex sequence* of D as a sequence of vertex sets of D say $(C_1, \ldots, C_t)$, such that,

1. for all $i, j \in [t]$, $i \neq j$, $C_i \cap C_j = \emptyset$ and $C_1 \uplus \ldots \uplus C_t = V(D)$, and
2. for all $i \in [t]$, $|C_i| \leq \gamma$,

For any subset $C_i \subseteq V(D)$, *deoct*(C_i) denotes the size of the minimum *deoct* of $D[C_i]$. The *cost* of a γ-vertex sequence $(C_1, \ldots, C_t)$ of D is defined as $\sum_{i\in[t]} deoct(C_i) + |\{(u, v) : (u, v) \in E(D), u \in C_j, v \in C_i, j > i\}|$. For the rest of the section, fix $\gamma = \alpha + \sqrt{\alpha^2 + 8\alpha k}$.

▶ **Lemma 8.3.** *Let $D \in \mathcal{D}_\alpha$. For any positive integer k, (D, k) is* YES *instance of* DEOCT *if and only if there exists a γ-vertex sequence of D of cost at most k.*

Proof. For the forward direction, let (D, k) be a YES instance of DEOCT. Let S be a *deoct* of D of size at most k. Consider the digraph $D - S$. Note that the vertex set of $D - S$ is the same as the vertex set of D. Let $(C_1, \ldots, C_t)$ be the topological ordering of the strongly connected components of $D - S$, that is, each C_i is a strongly connected component of $D - S$ and if there exists $(u, v) \in E(D - S)$, such that $u \in C_j$ and $v \in C_i$, then $j > i$. We will

now show that $(C_1, \ldots, C_t)$ is a γ-vertex sequence of D of cost at most k. First observe that $C_1 \uplus \ldots C_t = V(D)$. We will now show that for each $i \in [t]$, $|C_i| \leq \gamma$ and cost of $(C_1, \ldots, C_t)$ is at most k.

▷ Claim 8.4. For any $i \in [t]$, $|C_i| \leq \gamma$.

Proof. Since C_i is a strongly connected component of $D - S$ and S is a *deoct* of D, from Proposition 2.3 C_i is a bipartite graph in $D - S$. Let (A_i, B_i) be a bipartition of C_i. We will now show that $|A_i|, |B_i| \leq \frac{\gamma}{2}$. This will prove the claim. Let us argue that $|A_i| \leq \frac{\gamma}{2}$, the other case is symmetric. Since $D[A_i]$ is a subgraph of D, $D[A_i] \in \mathcal{D}_\alpha$. Thus, from Lemma 2.1, $E(D[A_i]) \geq \frac{|A_i|^2}{2\alpha} - \frac{|A_i|}{2}$. Since, A_i is an independent set in $D - S$, $|S| \geq E(D[A_i]) \geq \frac{|A_i|^2}{2\alpha} - \frac{|A_i|}{2}$. Then if $|A_i| > \frac{\gamma}{2}$, the we have that $|S| > k$, which is a contradiction. The same argument holds for B_i too. Thus, we conclude that $|C_i| \leq \gamma$. ◀

▷ Claim 8.5. The cost of $(C_1, \ldots, C_t)$ is at most k.

Proof. To show this, we will prove that the cost of $(C_1, \ldots, C_t)$ is at most $|S|$. Recall that cost of $(C_1, \ldots, C_t)$ is $\sum_{i \in [t]} \mathit{deoct}(C_i) + |\{(u, v) : (u, v) \in E(D), u \in C_j, v \in C_i, j > i\}|$. Let us denote $E_{back} = \{(u, v) : (u, v) \in E(D), u \in C_j, v \in C_i, j > i\}$. Since $(C_1, \ldots, C_t)$ is a topological ordering of the strongly connected components of $D - S$, $E_{back} \subseteq S$. Also, for any $i \in [t]$, $|\mathit{deoct}(C_i)| \leq |S \cap E(D[C_i])|$. Thus, cost of $(C_1, \ldots, C_t)$ is at most $|S|$. ◀

Claims 8.4 and 8.5 prove the forward direction of the lemma. We now prove the backward direction. Let $(C_1, \ldots, C_t)$ be a γ-vertex sequence of D of cost at most k. Let $E_{back} = \{(u, v) : (u, v) \in E(D), u \in C_j, v \in C_i, j > i\}$. We will now show that $S = \cup_{i \in [t]} \mathit{deoct}(C_i) \cup E_{back}$ is a *deoct* of D. Observe that $|S|$ is equal to the cost of $(C_1, \ldots, C_t)$. Suppose S is not a *deoct* of D. Then there exists an odd cycle in $D - S$. Since, for all $i \in [t]$, $\mathit{deoct}(C_i) \subseteq S$, such a cycle cannot be fully contained in any C_i. Therefore, there exists an arc of this cycle, say (u, v), such that $u \in C_j$ and $v \in C_i$, $j > i$. This violates that $E_{back} \subseteq S$. ◀

Let (D, k) be the input instance of DEOCT. The algorithm of Theorem 1.8 applies the kernelization algorithm of Theorem 1.6 to obtain an equivalent instance (D', k'). This is followed by a dynamic programming procedure over the k'-cuts in D' to obtain a γ-vertex sequence of D' of cost at most k.

Proof of Theorem 1.8. We will solve DEOCT by doing a dynamic programming over the set $\mathcal{C}$ of k-cuts. Let (D, k) be the input instance. Apply the kernelization algorithm of Theorem 1.6 to obtain an equivalent instance (D', k') where the number of vertices in D' is $k^{f(\alpha)}$. Since (D', k') is equivalent to (D, k), it is enough to solve the problem on (D', k'). For ease of notation, we will denote (D', k') by (D, k).

From Lemma 8.1, the number of k-cuts in D is at most $\eta = 2^{c'\sqrt{k}\log k} \cdot (k^{f(\alpha)} + 1)^{2\alpha\lceil\sqrt{k}\rceil} \cdot \log k$, where c' is a fixed absolute constant. From Lemma 8.2, all these k-cuts can be enumerated in $\eta \cdot k^{\mathcal{O}(1)}$ time. Let us denote by $\mathcal{C}$, the set of k-cuts of D.

Let T denote the dynamic programming table indexed by cuts in and integers $\{0, \ldots, k\}$. For any k-cut $(L, R) \in \mathcal{C}$ and $i \in \{0, \ldots, k\}$, we $T((L, R), i)$ is defined as follows.

$$T((L,R),i) = \begin{cases} 1 & \text{if there exists a } \gamma\text{-vertex sequence } (C_1, \ldots, C_\ell) \text{ of } D[L] \\ & \text{of cost at most } i, \text{ and } (L \setminus \{C_\ell\}, R \cup \{C_\ell\}) \in \mathcal{C} \\ 0 & \text{otherwise} \end{cases}$$

Note that $T((V(D),\emptyset),k) = 1$ if and only if D has a *deoct* of size at most k. This follows from Lemma 8.3.

We now describe how we compute $T((L,R),i)$, for any $(L,R) \in \mathcal{C}$ and $i \in [k]$. For all $i \in [k]$, $T((\emptyset,V(D)),i) = 1$. For any $(L,R) \in \mathcal{C}$, such that $L \neq \emptyset$, and any $i \in [k]$, $T((L,R),i) = 1$ if and only if the following holds: there exists $C \subseteq L$ such that $(L \setminus C, R \cup C) \in \mathcal{C}$, and $T((L \setminus C, R \cup C), i-j) = 1$ where $j = \mathit{deoct}(C) + |\{(u,v) : u \in C, v \in L \setminus C\}|$. Observe that the above describes a recursive procedure that computes all entries in T in time $2^{c^* f(\alpha)\sqrt{k}\log k}$ where c^* is an absolute constant. In total the running time of our algorithm is $2^{cf(\alpha)\sqrt{k}\log k} \cdot n^{\mathcal{O}(1)}$ where c is an absolute constant.

It only remains to prove the correctness of the above procedure. We now prove that for any $(L,R) \in \mathcal{C}$ and $i \in [k]$, $T((L,R),i) = 1$ if and only if there exists a γ-vertex sequence $(C_1,\ldots,C_\ell)$ of $D[L]$ of cost at most i, and $(L \setminus \{C_\ell\}, R \cup \{C_\ell\}) \in \mathcal{C}$. We prove this by induction on $|L|$. When $|L| = 0$, this is true because of the base case.

In the forward direction, we will show that if $T((L,R),i) = 1$ then there exists a γ-vertex sequence $(C_1,\ldots,C_\ell)$ of $D[L]$ of cost at most i, and $(L \setminus \{C_\ell\}, R \cup \{C_\ell\}) \in \mathcal{C}$. In the above procedure for computing the table T, we set $T((L,R),i) = 1$ only if there exists $C_\ell \subseteq L$, such that $(L \setminus \{C_\ell\}, R \cup \{C_\ell\}) \in \mathcal{C}$ and $T((L \setminus C, R \cup C), i-j) = 1$ where $j = \mathit{deoct}(C) + |\{(u,v) : u \in C, v \in L \setminus C\}|$. Since $T((L \setminus \{C_\ell\}, R \cup \{C_\ell\}), i-j) = 1$, by the induction hypothesis, $D[L \setminus \{C_\ell\}]$ has a γ-vertex sequence of cost at most $i-j$. Let $(C_1,\ldots,C_{\ell-1})$ be the ordering of $L \setminus \{C_\ell\}$ witnessing this, that is, cost of this ordering is at most $i-j$. Since $\mathit{deoct}(C) + |\{(u,v) : u \in C, v \in L \setminus C\}| = j$, the cost of $(C_1,\ldots,C_\ell)$ is at most i. Thus, the $(C_1,\ldots,C_{\ell-1},C_\ell)$ is a γ-vertex sequence of $D[L]$ of cost at most i.

In the reverse direction, we will show that if $D[L]$ has a γ-vertex sequence of cost at most i and $(L \setminus \{C_\ell\}, R \cup \{C_\ell\}) \in \mathcal{C}$, then $T((L,R),i) = 1$. Let $(C_1,\ldots,C_\ell)$ be a γ-vertex sequence in $D[L]$ of cost at most i such that $(L \setminus \{C_\ell\}, R \cup \{C_\ell\}) \in \mathcal{C}$. Let $j = \mathit{deoct}(C) + |\{(u,v) : u \in C, v \in L \setminus C\}|$. Then the sequence $(C_1,\ldots,C_{\ell-1})$ is a γ-vertex sequence of $D[L \setminus \{C_\ell\}]$ of cost at most $i-j$. Thus, $T((L \setminus \{C_\ell\}, R \cup \{C_\ell\}), i-j) = 1$. Then it follows that our recursive procedure sets $T((L,R),i) = 1$. This concludes the proof. ◀

9 Conclusion

In this paper, we presented a sparsification procedure for the class of acyclic digraphs (or more generally, "almost" acyclic) of bounded independence, to preserve the (both normal and parity-based) reachability from a given terminal set S to a given terminal set T under the failure of any set of at most k arcs. In particular, it outputs a digraph whose size is completely *independent* of n and polynomial in k, while even the simple classes of directed paths and tournaments admit no sparsifier whose output is a digraph of less than $n-1$ arcs already when $k = 1$. Apart from being interesting on its own from the perspective of fault tolerance, we also showed that our sparsification procedure finds applications in Kernelization. Specifically, we proved that the classic Directed Feedback Arc Set problem as well as Directed Edge Odd Cycle Transversal (which, in sharp contract, is W[1]-hard on general digraphs) admit polynomial kernels on bounded independence number digraphs. In fact, for any $p \in \mathbb{N}$, we designed a polynomial kernel for hitting all cycles of length ℓ where ($\ell \bmod p = 1$). Additionally, we derived complementary results that assert the NP-hardness of DEOCT on tournaments, as well as its admittance of a sub-exponential time parameterized algorithm on digraphs of bounded independence.

We conclude the paper with a few directions for further research. Our result, currently, holds when the input digraph D is "almost acyclic" and has bounded independence number.

From the example of the tournament described in the introduction (the one that is obtained by taking a transitive tournament and reversing the arcs along the Hamiltonian path defined by its topological ordering), it seems that some notion of "almost acyclic" might be necessary to have fault tolerant subgraphs whose size avoid the dependence on n. On the other hand, it might be possible to ask for something weaker than bounded independence number. For example, forbidding the existence of an induced P_α, the directed path on α vertices.

Question 1: Does $\mathrm{FTR}(S, S)$ admit a subgraph of size independent of n on digraphs that are "almost acyclic" and have no induced P_α, for some fixed positive integer α?

It is not very difficult to observe that our results (Lemmas 1.1 and 1.2) also hold when the input graph is undirected and has bounded independence number. It would be interesting (because of the arguments discussed earlier) if one could obtain similar results when the input undirected graph has no induced P_α.

Question 2: Does $\mathrm{FTR}(S, S)$ admit a subgraph of size independent of n when the input graph is undirected and has no induced P_α, for some fixed positive integer α?

It would also be interesting to discover other (di)graph classes where the dependence on n of the size of the output subgraph can be sublinear, for example, $\log n$, for $\mathrm{FTR}(S, S)$ and also for other fault tolerant graph properties.

References

1 Open problems in parameterized complexity. `http://fpt.wikidot.com/open-problems`. Accessed: 2019-02-15.

2 Nir Ailon, Moses Charikar, and Alantha Newman. Aggregating inconsistent information: Ranking and clustering. *J. ACM*, 55(5):23:1–23:27, 2008.

3 Noga Alon. Ranking tournaments. *SIAM J. Discrete Math.*, 20(1):137–142, 2006.

4 Noga Alon, Daniel Lokshtanov, and Saket Saurabh. Fast FAST. In *Automata, Languages and Programming, 36th International Colloquium, ICALP 2009, Rhodes, Greece, July 5-12, 2009, Proceedings, Part I*, pages 49–58, 2009.

5 Alexandr Andoni, Anupam Gupta, and Robert Krauthgamer. Towards (1+ ε)-approximate flow sparsifiers. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 279–293. Society for Industrial and Applied Mathematics, 2014.

6 Jørgen Bang-Jensen and Carsten Thomassen. A polynomial algorithm for the 2-path problem for semicomplete digraphs. *SIAM J. Discrete Math.*, 5(3):366–376, 1992.

7 Surender Baswana, Keerti Choudhary, Moazzam Hussain, and Liam Roditty. Approximate single source fault tolerant shortest path. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1901–1915. SIAM, 2018.

8 Surender Baswana, Keerti Choudhary, and Liam Roditty. Fault tolerant reachability for directed graphs. In *International Symposium on Distributed Computing*, pages 528–543. Springer, 2015.

9 Surender Baswana, Keerti Choudhary, and Liam Roditty. Fault-tolerant subgraph for single-source reachability: General and optimal. *SIAM Journal on Computing*, 47(1):80–95, 2018.

10 Stéphane Bessy, Fedor V Fomin, Serge Gaspers, Christophe Paul, Anthony Perez, Saket Saurabh, and Stéphan Thomassé. Kernels for feedback arc set in tournaments. *Journal of Computer and System Sciences*, 77(6):1071–1078, 2011.

11 Davide Bilò, Fabrizio Grandoni, Luciano Gualà, Stefano Leucci, and Guido Proietti. Improved purely additive fault-tolerant spanners. In *Algorithms-ESA 2015*, pages 167–178. Springer, 2015.

12 Davide Bilò, Luciano Gualà, Stefano Leucci, and Guido Proietti. Fault-tolerant approximate shortest-path trees. In *European Symposium on Algorithms*, pages 137–148. Springer, 2014.

13 Greg Bodwin. Linear size distance preservers. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 600–615. Society for Industrial and Applied Mathematics, 2017.

14 Diptarka Chakraborty and Debarati Das. Sparse weight tolerant subgraph for single source shortest path. In *16th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

15 Pierre Charbit, Stéphan Thomassé, and Anders Yeo. The minimum feedback arc set problem is np-hard for tournaments. *Combinatorics, Probability & Computing*, 16(1):1–4, 2007.

16 Moses Charikar, Tom Leighton, Shi Li, and Ankur Moitra. Vertex sparsifiers and abstract rounding algorithms. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 265–274. IEEE, 2010.

17 Shiri Chechik. Fault-tolerant compact routing schemes for general graphs. In *International Colloquium on Automata, Languages, and Programming*, pages 101–112. Springer, 2011.

18 Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. Fault tolerant spanners for general graphs. *SIAM Journal on Computing*, 39(7):3403–3423, 2010.

19 Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. F-sensitivity distance oracles and routing schemes. *Algorithmica*, 63(4):861–882, 2012.

20 Jianer Chen, Yang Liu, Songjian Lu, Barry O'sullivan, and Igor Razgon. A fixed-parameter algorithm for the directed feedback vertex set problem. *Journal of the ACM (JACM)*, 55(5):21, 2008.

21 Zhibin Chen, Jie Ma, and Wenan Zang. Coloring digraphs with forbidden cycles. *Journal of Combinatorial Theory, Series B*, 115:210 – 223, 2015. URL: `http://www.sciencedirect.com/science/article/pii/S0095895615000830`, `doi:https://doi.org/10.1016/j.jctb.2015.06.001`.

22 Rajesh Chitnis. *Directed Graphs: Fixed-Parameter Tractability & Beyond*. PhD thesis, 2014.

23 Rajesh Chitnis and Mohammad Taghi Hajiaghayi. Shadowless solutions for fixed-parameter tractability of directed graphs. *Encyclopedia of Algorithms*, pages 1–5, 2008.

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

25 Marek Cygan, Fedor V Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized algorithms*, volume 4. Springer, 2015.

26 Erik Demaine, Gregory Z Gutin, Dániel Marx, and Ulrike Stege. 07281 open problems–structure theory and fpt algorithmcs for graphs, digraphs and hypergraphs. In *Dagstuhl Seminar Proceedings*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2007.

27 Reinhard Diestel. Graph theory, volume 173 of. *Graduate texts in mathematics*, page 7, 2012.

28 Michael Dinitz and Robert Krauthgamer. Fault-tolerant spanners: better and simpler. In *Proceedings of the 30th annual ACM SIGACT-SIGOPS symposium on Principles of distributed computing*, pages 169–178. ACM, 2011.

29 Michael Dom, Jiong Guo, Falk Hüffner, Rolf Niedermeier, and Anke Truß. Fixed-parameter tractability results for feedback set problems in tournaments. In *Italian Conference on Algorithms and Complexity*, pages 320–331. Springer, 2006.

30 Rodney G Downey and Michael R Fellows. *Fundamentals of parameterized complexity*, volume 4. Springer, 2013.

31 Ran Duan and Seth Pettie. Connectivity oracles for graphs subject to vertex failures. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 490–509. SIAM, 2017.

32 Matthias Englert, Anupam Gupta, Robert Krauthgamer, Harald Racke, Inbal Talgam-Cohen, and Kunal Talwar. Vertex sparsifiers: New results from old techniques. *SIAM Journal on Computing*, 43(4):1239–1262, 2014.

33 Uriel Feige. Faster fast(feedback arc set in tournaments). *CoRR*, abs/0911.5094, 2009. URL: `http://arxiv.org/abs/0911.5094`.

34 Paola Festa, Panos M Pardalos, and Mauricio GC Resende. Feedback set problems. In *Handbook of combinatorial optimization*, pages 209–258. Springer, 1999.

35 Fedor V Fomin and Michał Pilipczuk. Subexponential parameterized algorithm for computing the cutwidth of a semi-complete digraph. In *European Symposium on Algorithms*, pages 505–516. Springer, 2013.

36 Alexandra Fradkin and Paul Seymour. Edge-disjoint paths in digraphs with bounded independence number. *Journal of Combinatorial Theory, Series B*, 110:19–46, 2015.

37 T. Gallai. On directed paths and circuits. *In: Erdös, P. and Katona, G., Eds., Theory of Graphs, Academic Press, New York*, pages 115–118, 1968.

38 Georges Gardarin and Wesley W. Chu. Integrity of databases: A general lockout algorithm with deadlock avoidance. In *Proceedings of the IFIP Working Conference on Modelling in Data Base Management Systems, [Freudenstadt, Germany, 5-8 January 1976]*, pages 395–411. North-Holland Publishing Company.

39 Richard M Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.

40 Marek Karpinski and Warren Schudy. Faster algorithms for feedback arc set tournament, kemeny rank aggregation and betweenness tournament. In *Algorithms and Computation - 21st International Symposium, ISAAC 2010, Jeju Island, Korea, December 15-17, 2010, Proceedings, Part I*, pages 3–14, 2010.

41 Claire Kenyon-Mathieu and Warren Schudy. How to rank with few errors. In *STOC*, volume 7, pages 95–103, 2007.

42 F Thomson Leighton and Ankur Moitra. Extensions and limits to vertex sparsification. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 47–56. ACM, 2010.

43 Daniel Lokshtanov, Misra Pranabendu, Saket Saurabh, and Meirav Zehavi. A brief note on single source fault tolerant reachability. *arXiv preprint arXiv:1904.08150*, 2019.

44 Daniel Lokshtanov, MS Ramanujan, and Saket Saurabh. When recursion is better than iteration: A linear-time algorithm for acyclicity with few error vertices. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1916–1933. SIAM, 2018.

45 Daniel Lokshtanov, MS Ramanujan, Saket Saurabh, and Meirav Zehavi. Parameterized complexity and approximability of directed odd cycle transversal. *arXiv preprint arXiv:1704.04249*, 2017.

46 Dániel Marx. Some open problems in parameterized complexity y. `http://www.cs.bme.hu/dmarx/papers/marx-dagstuhl2017-open.pdf`, 2-17.

47 Dániel Marx. What's next? future directions in parameterized complexity. In *The Multivariate Algorithmic Revolution and Beyond*, pages 469–496. Springer, 2012.

48 Pranabendu Misra, Saket Saurabh, Roohani Sharma, and Meirav Zehavi. Sub-exponential time parameterized algorithms for graph layout problems on digraphs with bounded independence number. In *38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

49 Hiroshi Nagamochi and Toshihide Ibaraki. A linear-time algorithm for finding a sparsek-connected spanning subgraph of ak-connected graph. *Algorithmica*, 7(1):583–596, Jun 1992. `doi:10.1007/BF01758778`.

50 Merav Parter. Dual failure resilient bfs structure. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing*, pages 481–490. ACM, 2015.

51 Merav Parter. Vertex fault tolerant additive spanners. *Distributed Computing*, 30(5):357–372, 2017.

52 Merav Parter et al. Fault-tolerant logical network structures. *Bulletin of EATCS*, 1(118), 2016.

53 Merav Parter and David Peleg. Sparse fault-tolerant bfs trees. In *European Symposium on Algorithms*, pages 779–790. Springer, 2013.

54 Mihai Patrascu and Mikkel Thorup. Planning for fast connectivity updates. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 263–271. IEEE, 2007.

55 Venkatesh Raman and Saket Saurabh. Parameterized algorithms for feedback set problems and their duals in tournaments. *Theoretical Computer Science*, 351(3):446–458, 2006.

56 B. Roy. Nombre chromatique et plus longs chemins d'un graphe. *Rev. Francaise Informat. Recherche Operationnelle*, 1:129–132, 1967.

57 Abraham Silberschatz, Greg Gagne, and Peter B Galvin. *Operating system concepts (9th edition)*. Wiley, 2012.

58 Ewald Speckenmeyer. On feedback problems in digraphs. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 218–231. Springer, 1989.