

# Parameterized Complexity and Approximability of Directed Odd Cycle Transversal\*

Daniel Lokshтанov<sup>†</sup>   M. S. Ramanujan<sup>‡</sup>   Saket Saurabh<sup>§¶</sup>   Meirav Zehavi<sup>||</sup>

## Abstract

A *directed odd cycle transversal* of a directed graph (digraph)  $D$  is a vertex set  $S$  that intersects every *odd directed cycle* of  $D$ . In the DIRECTED ODD CYCLE TRANSVERSAL (DOCT) problem, the input consists of a digraph  $D$  and an integer  $k$ . The objective is to determine whether there exists a directed odd cycle transversal of  $D$  of size at most  $k$ . In this paper, we settle the parameterized complexity of DOCT when parameterized by the solution size  $k$  by showing that DOCT does not admit an algorithm with running time  $f(k)n^{\mathcal{O}(1)}$  unless  $\text{FPT} = \text{W}[1]$ . On the positive side, we give a factor 2 fixed-parameter approximation (FPT approximation) algorithm for the problem. More precisely, our algorithm takes as input  $D$  and  $k$ , runs in time  $2^{\mathcal{O}(k^2)}n^{\mathcal{O}(1)}$ , and either concludes that  $D$  does not have a directed odd cycle transversal of size at most  $k$ , or produces a solution of size at most  $2k$ . Finally, assuming  $\text{gap-ETH}$ , we show that there exists an  $\epsilon > 0$  such that DOCT does not admit a factor  $(1 + \epsilon)$  FPT-approximation algorithm. In fact, our lower bound holds even under the weaker hypothesis that for some  $\epsilon > 0$ , there is no  $f(k)n^{\mathcal{O}(1)}$  time algorithm that distinguishes between a satisfiable Binary CSP instance from one where every assignment violates at least an  $\epsilon$  fraction of the constraints.

---

\*Supported by *Pareto-Optimal Parameterized Algorithms*, ERC Starting Grant 715744 and *Parameterized Approximation*, ERC Starting Grant 306992. M. S. Ramanujan also acknowledges support from *BeHard*, Bergen Research Foundation and *X-Tract*, Austrian Science Fund (FWF, project P26696).

<sup>†</sup>Department of Computer Science, University of California, Santa Barbara, USA. [daniello@ucsb.edu](mailto:daniello@ucsb.edu)

<sup>‡</sup>University of Warwick, Warwick, United Kingdom. [R.Maadapuzhi-Sridharan@warwick.ac.uk](mailto:R.Maadapuzhi-Sridharan@warwick.ac.uk)

<sup>§</sup>The Institute of Mathematical Sciences, HBNI, Chennai, India. [saket@imsc.res.in](mailto:saket@imsc.res.in)

<sup>¶</sup>University of Bergen, Norway

<sup>||</sup>Ben-Gurion University, Beersheba, Israel. [meiravze@bgu.ac.il](mailto:meiravze@bgu.ac.il)

# 1 Introduction

A *directed odd cycle transversal* of a digraph  $D$  is a set  $S$  of vertices of  $D$  such that deleting  $S$  from  $D$  results in a graph without any directed odd cycles. In the NP-complete DIRECTED ODD CYCLE TRANSVERSAL (DOCT) problem, the input consists of a digraph  $D$  on  $n$  vertices and an integer  $k$ , and the task is to determine whether  $D$  has a directed odd cycle transversal of size at most  $k$ . DOCT generalizes several well studied problems such as ODD CYCLE TRANSVERSAL (OCT) on undirected graphs [1, 15, 31, 48], DIRECTED FEEDBACK VERTEX SET (DFVS) [8, 22, 24, 30], and DIRECTED SUBSET FEEDBACK VERTEX SET [13, 22]. In OCT, the input consists of an undirected graph  $G$  and integer  $k$ , and the task is to determine whether there exists a subset  $S$  of vertices such that  $G - S$  is bipartite.<sup>1</sup> In DFVS, the input consists of a digraph  $D$  and integer  $k$ , and the task is to determine whether there exists a subset  $S$  of vertices such that  $D - S$  is a directed acyclic graph.<sup>2</sup>

The existence of fixed-parameter algorithms (FPT algorithms) for OCT and DFVS were considered to be major open problems in parameterized complexity, until FPT algorithms were found for OCT in 2003 by Reed et al. [48], and for DFVS in 2007 by Chen et al. [8]. The algorithms for these two problems have had significant influence on the development of the field, resulting in proliferation of techniques such as *iterative compression* and *important separators* [16, 21]. Today both algorithms and corresponding methods are considered fundamental textbook material [16].

Once both OCT and DFVS were shown to be FPT, DOCT immediately became the next natural target. The parameterized complexity of DOCT was explicitly stated as an open problem [18] for the first time in 2007, immediately after the announcement of an FPT algorithm for DFVS. Since then the problem has been re-stated several times [9, 12, 42, 43]. In this paper, we settle the parameterized complexity of DOCT, by showing that the problem is W[1]-hard. Our hardness proof also gives a near-tight running time lower bound for DOCT assuming the Exponential Time Hypothesis (ETH). In particular, we prove the following.

**Theorem 1.** *DOCT is W[1]-hard. Furthermore, assuming the ETH there is no algorithm for DOCT with running time  $f(k)n^{o(k/\log k)}$ .*

On the one hand, Theorem 1 shows that DOCT is intractable from the perspective of parameterized complexity. On the other hand, the problem is known not to admit a constant factor approximation algorithm running in polynomial time, assuming the Unique Games Conjecture [31]. Hence, the next natural question is whether one could get a constant factor approximation algorithm in FPT time. Our second result is an affirmative answer to this question.

**Theorem 2.** *DOCT admits a  $2^{\mathcal{O}(k^2)}n^{\mathcal{O}(1)}$  time FPT-approximation algorithm with approximation ratio 2.*

In fact, Theorem 2 follows as a corollary from a stronger result for a “labeled digraph problem” that we introduce. We show that this problem subsumes DOCT as well as the NODE UNIQUE LABEL COVER problem [10, 28, 37], and design an FPT approximation algorithm that works even for this more general problem.

In light of Theorem 2 the next natural question is whether the approximation factor can be made arbitrarily close to 1. Our final contribution is to provide evidence that there exists an  $\epsilon > 0$  such that DOCT does not admit a  $(1 + \epsilon)$  FPT-approximation algorithm. In particular, the proof of Theorem 1 can be thought of as a parameterized reduction from the BINARY CONSTRAINT SATISFACTION (BCSP) problem, informally defined as follows. The input consists

---

<sup>1</sup>OCT reduces to DOCT by replacing every edge by two arcs, one in each direction.

<sup>2</sup>DFVS reduces to DOCT by adding for every arc  $uv$  of  $D$  a new vertex  $x$  as well as the arcs  $ux$  and  $xv$ .

of two integers  $n$  and  $k$  specifying that there are  $k$  variables,  $x_1, \dots, x_k$ , each variable  $x_i$  taking a value from  $\{1, \dots, n\}$ , together with a list of *constraints*. Each constraint specifies two variables,  $x_i$  and  $x_j$ , together with a list  $L$  of all legal pairs of values that  $x_i$  and  $x_j$  may take simultaneously. An assignment of values to the variables satisfies the constraint if  $(x_i, x_j) \in L$ . The task is to find an assignment that satisfies all constraints. It is well known (see e.g. [41]) that BCSP parameterized by the number of variables  $k$  is W[1]-complete. We conjecture that not only is it W[1]-hard to find a satisfying assignment to a BCSP instance if there is one, but it is also W[1]-hard to distinguish between instances that have a satisfying assignment from instances where every assignment violates at least an  $\epsilon$  fraction of the constraints. Formally, for every  $\epsilon > 0$ , we define the promise problem  $\epsilon$ -GAP-BCSP, as BCSP where the input instance is promised to either be satisfiable, or have the property that every assignment violates at least an  $\epsilon$  fraction of the constraints. The task is to determine whether the input instance is satisfiable or not.

**Hypothesis 1 (Parameterized Inapproximability Hypothesis (PIH)).** *There exists an  $\epsilon > 0$  such that  $\epsilon$ -GAP-BCSP is W[1]-hard.*

PIH is strongly related to the recently introduced Gap-Exponential Time Hypothesis (Gap-ETH) [20, 39], which is a stronger version of the well-known Exponential Time Hypothesis [25, 26]. Moreover, Gap-ETH has been instrumental in establishing several parameterized inapproximability results in recent years [6, 5, 11].

We remark that for purposes of showing hardness of approximation results starting from Gap-ETH, we could just as well have conjectured that there exists an  $\epsilon > 0$  such that there is no  $f(k)n^{\mathcal{O}(1)}$  time algorithm for  $\epsilon$ -GAP-BCSP. This is because (see Section 4.2 for a simple proof) assuming Gap-ETH, there exists an  $\epsilon > 0$  such that there is no  $f(k)n^{\mathcal{O}(1)}$  time algorithm for  $\epsilon$ -GAP-BCSP<sup>3</sup>. Thus, (essentially) any hardness of approximation result assuming PIH can be shown assuming Gap-ETH instead. In addition, the recent result of Bhattacharya et al. [5] on the parameterized inapproximability of the famous EVEN SET problem is based on this assumption which is slightly weaker than that in the statement of PIH above. However, we strongly believe that PIH is true as stated—indeed, we should hardly claim this conjecture as our own, as quite a few researchers in parameterized complexity have stated this conjecture as a natural formulation of a PCP-theorem in the context of parameterized inapproximability<sup>4</sup>.

Our final result is that assuming either PIH or Gap-ETH, there exists an  $\epsilon > 0$  such that DOCT does not admit an FPT-approximation algorithm with ratio  $1 + \epsilon$ .

**Theorem 3.** *Assuming Gap-ETH or PIH and  $\text{FPT} \neq \text{W}[1]$ , there exists an  $\epsilon > 0$  such that DOCT does not admit an FPT-approximation algorithm with approximation ratio  $1 + \epsilon$ .*

**Arc-Directed Odd Cycle Transversal.** We remark that easy reductions transfer all of our results to ARC-DOCT, the “arc” version of DOCT where the goal is to remove at most  $k$  arcs such that the resulting graph does not have any directed odd cycles. To transfer the hardness results we need to reduce DOCT to ARC-DOCT. For this purpose, it is sufficient to subdivide every arc, and then split every original vertex  $u$  of the input digraph into two vertices,  $u_{in}$  and  $u_{out}$ , such that all arcs leading into  $u$  lead into  $u_{in}$  instead, all arcs leading out of  $u$  lead out of  $u_{out}$  instead, and adding the arc  $u_{in}u_{out}$ . To transfer the algorithmic results from DOCT to ARC-DOCT, we need to reduce ARC-DOCT to DOCT. This is achieved by subdividing every arc twice, and then making each original vertex undeletable by adding  $k + 1$  copies of it.

<sup>3</sup>Bhattacharya et al. have also included a proof of this statement in [4].

<sup>4</sup>Such statements have always been made to us as private communications by various researchers in informal settings. Therefore we do not include a citation for this claim.

## Our Methods

**W[1]-hardness.** The starting point for both our hardness results as well as our approximation algorithm is a failed attempt at obtaining an FPT algorithm. The root of this attempt was the FPT algorithm for DFVS by Chen et al. [8] (see also the textbook [16] for a more gentle exposition). The key concept in this algorithm is the notion of *important separators*, defined by Marx [40]. Given a digraph  $D$  and two vertices  $u$  and  $v$ , a  $u$ - $v$ -separator is a vertex set  $S \subseteq V(D) \setminus \{u, v\}$  such that there is no directed path from  $u$  to  $v$  in  $D - S$ . A  $u$ - $v$ -separator  $S$  is called a *minimal  $u$ - $v$ -separator* if no proper subset of  $S$  is also a  $u$ - $v$ -separator.

Given a vertex set  $S$  such that  $u$  is not in  $S$ , we define the *reach of  $u$  in  $D - S$*  as the set  $R_D(u, S)$  of vertices reachable from  $u$  by a directed path in  $D - S$ . We can now define a partial order on the set of minimal  $u$ - $v$  separators as follows. Given two minimal  $u$ - $v$  separators  $S_1$  and  $S_2$ , we say that  $S_1$  is “at least as good as”  $S_2$  if  $|S_1| \leq |S_2|$  and  $R_D(u, S_2) \subseteq R_D(u, S_1)$ . In plain words,  $S_1$  “costs less” than  $S_2$  in terms of the number of vertices deleted, and  $S_1$  “is pushed further towards  $v$ ” than  $S_2$  is. A minimal  $u$ - $v$  separator  $S$  is an *important  $u$ - $v$ -separator* if no minimal  $u$ - $v$ -separators other than  $S$  is at least as good as  $S$ . The key insight behind the algorithm for DFVS by Chen et al. [8], as well as algorithms for several other parameterized problems [14, 13, 17, 32, 33, 36, 38, 37, 44], is that for every  $k$ , the number of important  $u$ - $v$ -separators of size at most  $k$  is at most  $4^k$  [7]. We refer the reader to the textbook by Cygan et al. [16] for a more thorough exposition of important separators.

Applying the initial steps of the DFVS algorithm to DOCT (i.e. the methods of iterative compression, and guessing an order on an undeletable solution), one naturally arrives at an extension of the notion of important separators. Let us define the *cleaning cost* of a minimal  $u$ - $v$  separator  $S$  as  $\text{doct}(D[R_D(u, S)])$ , where  $\text{doct}(D)$  is the minimum size of a directed odd cycle transversal of  $D$ . Then, we define a new partial order on minimal  $u$ - $v$  separators. Here, given two minimal  $u$ - $v$  separators  $S_1$  and  $S_2$ , we say that  $S_1$  is “at least as good as”  $S_2$  if  $|S_1| \leq |S_2|$ ,  $R_D(u, S_2) \subseteq R_D(u, S_1)$ , and the cleaning cost of  $S_1$  is at most the cleaning cost of  $S_2$ . In other words,  $S_1$  costs less than  $S_2$ ,  $S_1$  is pushed further towards  $v$  than  $S_2$ , and “cleaning up” the reach of  $u$  in  $G - S_1$  does not cost more than cleaning up the reach of  $u$  in  $G - S_2$ . We say that a minimal  $u$ - $v$  separator  $S$  is a *DOCT-important  $u$ - $v$ -separator* if no minimal  $u$ - $v$ -separators other than  $S$  are at least as good as  $S$  with respect to this new partial order.

For every digraph  $D$ , vertices  $u$  and  $v$  and integer  $k$ , we know that there are at most  $4^k$  important  $u$ - $v$  separators of size at most  $k$ . For the purposes of an FPT algorithm for DOCT, the pivotal question becomes whether the number of DOCT-important  $u$ - $v$ -separators of size at most  $k_1$  and cleaning cost at most  $k_2$  can be upper bounded by a function of  $k_1$  and  $k_2$  only, or if there exist families of graphs where the number of DOCT-important  $u$ - $v$ -separators of size at most  $k_1$  and cleaning cost at most  $k_2$  grows with the size of the graphs. Indeed, a constructive upper bound on  $f(k_1, k_2)$ , the number of DOCT-important  $u$ - $v$ -separators of size at most  $k_1$  and cleaning cost at most  $k_2$ , would have implied an FPT algorithm for DOCT.

We managed to prove that there exists a function  $f$  such that the number of DOCT-important  $u$ - $v$ -separators of size at most  $k$  and cleaning cost 0 is at most  $f(k)$ . In a subsequent attempt to similarly upper bound the number of DOCT-important  $u$ - $v$ -separators of size at most  $k$  and cleaning cost 1, we discovered the *clock gadgets* (see Section 3.2), which are graphs where the number of DOCT-important  $u$ - $v$ -separators of size at most 2 and cleaning cost 1 is  $\Omega(n)$ .

A clock gadget essentially permits us to encode (in the language of DOCT) the choice of one element out of a domain of size  $n$ , without it being clear a priori which element(s) should be the best one(s) to select. For many problems, once one has such a selection gadget it is easy to prove W[1]-hardness by reducing from BCSP (or, equivalently, from MULTICOLORED CLIQUE). However, we were able to show that on graphs consisting only of clock gadgets glued

together in the most natural way, DOCT is in fact FPT<sup>5</sup>. In particular, clocks do not provide a general way of synchronizing the choices of different elements, making it difficult to encode the constraints of BCSP using DOCT. We were able to engineer such a synchronization gadget by a non-trivial modification of the “grid gadget” used by Pilipeczuk and Wahlström [46] to show W[1]-hardness of DIRECTED MULTICUT with four terminal pairs. At this point one can complete a reduction from BCSP using clocks to encode the selection of a value for each variable and using synchronization gadgets to encode the constraints of the BCSP instance.

**FPT-Approximation.** The hardness of DOCT comes from the fact that DOCT-important  $u$ - $v$ -separators have to do two jobs at the same time. First, they need to disconnect  $v$  from  $u$ , and second they need to clean the reach of  $u$  from directed odd cycles. Our approximation algorithm works by delegating the two jobs to different solutions, and solving each of the jobs separately and optimally.

Just like our W[1]-hardness proof, our FPT-approximation for DOCT builds on the algorithm of Chen et al. [8] for DFVS. The method of iterative compression (see [16, 21]) allow us to reduce the original problem to the setting where we are given a digraph  $D$ , an integer  $k$ , and a directed odd cycle transversal  $\hat{S}$  of size  $2k + 1$ . The task is to either determine that  $D$  does not have a directed odd cycle transversal of size at most  $k$ , or output a directed odd cycle transversal of size at most  $2k$ . We now proceed with a sketch of how to solve this task in FPT time.

In order to witness that a digraph  $D$  has no directed odd cycles it is sufficient to partition the vertex set of  $D$  into sets  $Z_1, Z_2, \dots, Z_\ell$  such that (a) no arc goes from  $Z_i$  to  $Z_j$  with  $j < i$  and (b) for every  $i \leq \ell$  the underlying undirected graph of  $D[Z_i]$  is bipartite. To certify (b) it is sufficient to provide a coloring of all vertices in  $D$  with black or white, such that every arc with both endpoints in  $Z_i$  for some  $i$  has different colored endpoints. The sets  $Z_1, Z_2, \dots, Z_\ell$  can always be chosen to be the strongly connected components of  $D$ , and in this case the ordering  $Z_1, Z_2, \dots, Z_\ell$  can be any topological ordering of the directed acyclic graph obtained from  $D$  by collapsing every strongly connected component to a vertex.

Suppose now that  $D$  has a directed odd cycle transversal  $S$  of size at most  $k$ . Let  $Z_1, Z_2, \dots, Z_\ell$  be a partitioning of  $V(D - S)$  and  $\phi : V(D - S) \rightarrow \{\text{black}, \text{white}\}$  be a coloring that certifies that  $D - S$  does not have directed odd cycles. At the cost of a  $3^k$  overhead in the running time we can guess for each vertex  $v \in \hat{S}$  whether it is deleted (i.e put in the directed odd cycle transversal), colored black or colored white. At the cost of an additional  $k!$  overhead in the running time we can guess for every pair of vertices  $u, v$  in  $\hat{S}$  whether they occur in the same strong component  $Z_i$ , and if not, which of the two strong components containing  $u$  and  $v$  respectively occurs first in the ordering  $Z_1, Z_2, \dots, Z_\ell$ . Applying these guesses together with some simple reduction rules, we end up in the following setting. The input is a digraph  $D$ , an integer  $k$  and a set  $\hat{S}$  such that  $D - S$  contains no directed odd cycles, and  $D[S]$  is an acyclic tournament (that is, there is an arc between every pair of vertices in  $S$ ). The task is to either find a set  $S \subseteq V(D) \setminus \hat{S}$  of size at most  $2k$  such that **(a)**  $S$  is a directed odd cycle transversal, and **(b)** no strong component of  $D - S$  contains more than one vertex of  $\hat{S}$ , or to conclude that no such set of size at most  $k$  exists.

A set  $S$  that only satisfies **(b)** is called a *skew separator* for  $\hat{S}$ , and the main subroutine in the algorithm of Chen et al. [8] for DFVS is an algorithm that given  $D$ ,  $\hat{S}$  and  $k$ , runs in time  $\mathcal{O}(4^k k^{\mathcal{O}(1)}(n + m))$ , and finds a skew separator  $S$  for  $\hat{S}$  of size at most  $k$  if such a skew separator exists. Our approximation algorithm runs this subroutine and either finds a skew separator  $S$  of size at most  $k$ , or concludes that no set of size at most  $k$  can satisfy both **(a)** and **(b)** (in particular, just **(b)**). It then determines in time  $2^{\mathcal{O}(k^2)} n^{\mathcal{O}(1)}$  whether  $D - S$  has a directed odd cycle transversal of size at most  $k$  disjoint from  $\hat{S}$ . If such a set  $S^*$  exists, the algorithm outputs

<sup>5</sup>Because this is such a specialized graph class, we did not include a proof of this fact in the paper.

$S \cup S^*$  as a solution of size at most  $2k$  that satisfies **(a)** and **(b)**. If no such directed odd cycle transversal  $S^*$  exists, the approximation algorithm concludes that no set of size at most  $k$  can satisfy both **(a)** and **(b)** (in particular, just **(a)**). All that remains is to describe the algorithm for finding in time  $2^{\mathcal{O}(k^2)}n^{\mathcal{O}(1)}$  a directed odd cycle transversal  $S^*$  of size at most  $k$  disjoint from  $\hat{S}$  in  $D - S$ , or determining that such a set does not exist.

At this point we observe that the problem breaks up into independent sub-problems for each strongly connected component of  $D - S$ . For each such component  $C$  we have that  $|C \cap \hat{S}| \leq 1$ , because  $S$  is a skew separator for  $\hat{S}$ . Since  $\hat{S}$  is a directed odd cycle transversal for  $D$ , if  $C \cap \hat{S} = \emptyset$  then there can be no directed odd cycles in  $D[C]$ . Hence we concentrate on the case when  $C \cap \hat{S} = \{w\}$  for a vertex  $w$ . In other words, we are down to the case where the input is a digraph  $D$ , integer  $k$  and a vertex  $w$  such that  $\{w\}$  is a directed odd cycle transversal for  $D$ . The task is to find a directed odd cycle transversal  $S^*$  of  $D$  of size at most  $k$  with  $w \notin S^*$ .

Define the *shadow* of  $S^*$  to be the set of all vertices of  $D - S^*$  that are not in the strongly connected component of  $G - S^*$  containing  $w$ . Using the technique of *shadow removal*, introduced by Marx and Razgon [44] (see also [14, 13, 12]) in their FPT algorithm for MULTICUT, we can reduce the problem to the special case where the shadow of  $S^*$  is empty, at the cost of a  $2^{\mathcal{O}(k^2)}n^{\mathcal{O}(1)}$  overhead in the running time. In this special case  $D - S^*$  is strongly connected, and therefore the underlying undirected graph of  $D - S^*$  is bipartite. Thus,  $S^*$  is an *undirected* odd cycle transversal for the underlying undirected graph of  $D$ . Here we can apply any one of the numerous FPT algorithms [27, 34, 47, 48] for OCT. Thus we can find optimal directed odd cycle transversals in FPT time for the case when a single undeletable vertex is a directed odd cycle transversal, and as discussed above, this is sufficient to complete the factor 2 FPT-approximation.

As a subroutine of our FPT-approximation we gave an FPT algorithm for DOCT for the special case where an undeletable directed odd cycle transversal of size 1 is given as input. Our hardness result also holds for the case where an undeletable directed odd cycle transversal of size 3 is given as input (the vertices  $\{x, y, z\}$  in the construction). Therefore, the parameterized complexity of the case when one also has an undeletable directed odd cycle transversal of size 2 in the input, is an interesting open problem. It is conceivable that an FPT algorithm for this case could help in obtaining an FPT-approximation for DOCT with a factor better than 2.

We remark that this high-level approach extends to a more general problem that subsumes DOCT as well as NODE UNIQUE LABEL COVER. Therefore, we design our FPT approximation for the general problem and derive the algorithm for DOCT as a corollary.

**FPT-Inapproximability.** For every  $\epsilon > 0$  there exists a  $\delta > 0$  such that our first reduction, which proves the W[1]-hardness of DOCT, also translates the hardness of  $\epsilon$ -GAP-BCSP into hardness of distinguishing between digraphs  $D$  such that  $\text{doct}(D) \leq k$  from digraphs  $D$  such that  $\text{doct}(D) > k(1 + \delta)$ . However, the reduction only works if in the instance of  $\epsilon$ -GAP-BCSP every variable occurs in at most three constraints. To complete the proof of the parameterized inapproximability of DOCT, we need to reduce  $\epsilon$ -GAP-BCSP to this special case. We achieve this by replacing every “high degree” variable by a group of independent low degree variables, while ensuring that the low degree variables all get the same value by introducing a constant degree expander of equality constraints between them.

## 2 Preliminaries

We use the notations  $[t]$  and  $[t]_0$  as shorthands of  $\{1, 2, \dots, t\}$  and  $\{0, 1, \dots, t\}$ , respectively. Given a function  $f : A \rightarrow \mathbb{R}$  and a subset  $A' \subseteq A$ , denote  $f(A') = \sum_{a \in A'} f(a)$ .

**Parameterized Complexity.** Formally, a *parameterization* of a problem is the assignment of an integer  $k$  to each input instance. Here, the goal is to confine the combinatorial explosion in

the running time of an algorithm for  $\Pi$  to depend only on  $k$ . We say that a parameterized problem  $\Pi$  is *fixed-parameter tractable* (FPT) if there exists an algorithm that solves  $\Pi$  in time  $f(k) \cdot |I|^{\mathcal{O}(1)}$ , where  $|I|$  is the size of the input instance and  $f$  is an arbitrary computable function depending only on the parameter  $k$ .

On the negative side, parameterized complexity also provides methods to show that a problem is unlikely to be FPT. The main technique is the one of parameterized reductions analogous to those employed in classical complexity. Here, the concept of  $W[1]$ -hardness replaces the one of NP-hardness, and we need not only construct an equivalent instance in FPT time, but also ensure that the size of the parameter in the new instance depends only on the size of the parameter in the original instance. For our purposes, it is sufficient to note that if there exists such a reduction transforming a problem known to be  $W[1]$ -hard to another problem  $\Pi$ , then the problem  $\Pi$  is  $W[1]$ -hard as well. Central  $W[1]$ -hard-problems include, for example, the problem of deciding whether a nondeterministic single-tape Turing machine accepts within  $k$  steps, the CLIQUE problem parameterized by solution size, and the INDEPENDENT SET problem parameterized by solution size

In the context of a parameterized minimization problem  $\Pi$ , we say that an algorithm for  $\Pi$  is an  $\alpha$ -*approximation algorithm* if it always outputs a solution of size at most  $\alpha k$  when there exists a solution of size at most  $k$  (in other words, the input instance is a yes-instance), and it always outputs NO when there does not exist a solution of size at most  $\alpha k$ . Additional details can be found in the monographs [23, 45, 21, 16].

**Digraphs.** We refer to standard terminology from the book of Diestel [19] for those graph-related terms that are not explicitly defined here. Given a digraph  $D$  and a vertex set  $X \subseteq V(D)$ , we say that  $X$  is a *directed odd cycle transversal* of  $D$  if  $X$  intersects every directed odd cycle of  $D$ . We further say that  $X$  is a *minimal* directed odd cycle transversal of  $D$  if no proper subset of  $D$  is also a directed odd cycle transversal of  $D$ . Finally, we call  $X$  a *minimum* directed odd cycle transversal of  $D$  if there is no directed odd cycle transversal of  $D$  whose size is strictly smaller than the size of  $X$ . In the context of DOCT, we use the terms *solution* and  $\alpha$ -*approximate solution* to refer to directed odd cycle transversals of sizes at most  $k$  and at most  $\alpha k$ , respectively.

Given a vertex set  $X \subseteq V(D)$ , we let  $D[X]$  denote the subgraph of  $D$  induced by  $X$ , and we define  $D \setminus X = D[V(D) \setminus X]$ . Given an arc  $(u, v) \in A(D)$ , we refer to  $u$  as the *tail* of the arc and to  $v$  as the *head* of the arc. Given a vertex set  $X \subseteq V(G)$ , we use  $N^+(X)$  to denote the set of out-neighbors of  $X$  and  $N^-(X)$  to denote the set of in-neighbors of  $X$ . We use  $N^i[X]$  to denote the set  $X \cup N^i(X)$  where  $i \in \{+, -\}$ . We denote by  $A[X]$  the subset of edges in  $A(D)$  with both endpoints in  $X$ . A *strongly connected component* of  $D$  is a maximal subgraph in which every vertex has a directed path to every other vertex. We say that a strongly connected component is *non-trivial* if it consists of at least two vertices and *trivial* otherwise. For disjoint vertex sets  $X$  and  $Y$ , the set  $Y$  is said to be *reachable from  $X$*  if for *every* vertex  $y \in Y$ , there exists a vertex  $x \in X$  such that the  $D$  contains a directed path from  $x$  to  $y$ . For a vertex  $v \in V(D)$  and walk  $W = v_1, \dots, v_r$ , we say that  $W$  is a  *$v$ -walk* if there is an  $i \in [r]$  such that  $v = v_i$ . We say that  $W$  is a *closed  $v$ -walk* if  $v_1 = v_r = v$ . We say that  $W$  is an  *$x$ - $y$  walk* if  $v_1 = x$  and  $v_r = y$ . Sometimes we say that a  $v$ -walk is a  *$v$ - $v$  walk*. This is simply so that we can refer to  $x$ - $y$  walks in general without having to resort to a separate proof (when it is not necessary) for the case when  $x=y$ . For  $1 \leq i < j \leq r$ , we denote by  $W[v_i, v_j]$  the subwalk of  $W$  from  $v_i$  to  $v_j$ . We call the vertices  $v_2, \dots, v_{r-1}$ , the *internal* vertices of the walk  $W$ . For two walks  $W_1 = v_1, \dots, v_t$  and  $W_2 = w_1, \dots, w_q$  such that  $v_t = w_1$ , we denote by  $W_1 + W_2$  the concatenated walk  $v_1, \dots, v_{t-1}, v_t, w_2, \dots, w_q$ . For disjoint subsets  $X, Y, Z \subseteq V(D)$ , we call  $Z$  an  *$X$ - $Y$  separator* if there is no path from a vertex of  $X$  to a vertex of  $Y$  in  $D - Z$ .

Our proofs rely on the following well-known proposition (see, e.g., [3]).

**Proposition 2.1** (Folklore). *Let  $D$  be a strongly connected directed graph that does not contain*

a directed odd cycle. Then, the underlying undirected graph of  $D$  is a bipartite graph.

### 3 W[1]-Hardness

In this section, we resolve the question of the parameterized complexity of DOCT. More precisely, we prove Theorem 1. For convenience, let us restate the theorem below.

**Theorem 1.** DOCT is W[1]-hard. Furthermore, assuming the ETH there is no algorithm for DOCT with running time  $f(k)n^{o(k/\log k)}$ .

The source of our reduction is the PARTITIONED SUBGRAPH ISOMORPHISM (PSI) problem. The definition of this problem relies on the notion of a *colorful mapping*. Given undirected graphs  $H$  and  $G$  where  $G$  has maximum degree 3, and a coloring function  $col : V(H) \rightarrow V(G)$ , we say that an injective function  $\varphi : V(G') \rightarrow V(H)$  is a *colorful mapping of  $G'$  into  $H$* , where  $G'$  is a subgraph of  $G$ , if for every  $v \in V(G')$ ,  $col(\varphi(v)) = v$ , and for every  $\{u, v\} \in E(G')$ ,  $\{\varphi(u), \varphi(v)\} \in E(H)$ . Formally, the PSI problem is defined as follows.

PARTITIONED SUBGRAPH ISOMORPHISM (PSI)	
<i>Input:</i>	Undirected graphs $H$ and $G$ , and a coloring function $col : V(H) \rightarrow V(G)$ . The maximum degree of a vertex of $G$ is 3.
<i>Question:</i>	Does there exist a colorful mapping of $G$ into $H$ ?

While the PSI problem requires us to map the entire graph  $G$ , to prove our inapproximability result we would also be interested in colorful mappings of *subgraphs* of  $G$ . In the context of the PSI problem, we rely on a well-known proposition due to Marx [41], where the same problem is called COLORED SUBGRAPH ISOMORPHISM.

**Proposition 3.1** (Corollary 6.3, [41]). *The PSI problem is W[1]-hard. Moreover, unless ETH fails, PSI cannot be solved in time  $f(k)n^{o(\frac{k}{\log k})}$  for any function  $f$  where  $k = |E(G)|$ . Here,  $n = |V(H)|$ .*

We anticipate that the components introduced by our proof will be useful for other reductions that aim to establish the W[1]-hardness of problems involving parities and/or cuts. Hence, we have structured our proof as follows. First, for the sake of clarity of the proof, we integrate arc and vertex annotations into the definition of DOCT. Then, we introduce the concept of a *clock*, which is a gadget that lies at the heart of our reduction. This gadget captures the power of parities in a compact, easy-to-use manner. In particular, it elegantly encodes the selection of two (not necessarily distinct) indices from a set  $[n]$  whose sum is upper bounded by  $n + 1$  (in the case of a *forward clock*) or lower bounded by  $n + 1$  (in the case of a *reverse clock*). We remark that the selection is orchestrated by a variable that we call *time*. Next, we “glue” the tips of the *hands* of a forward clock and a reverse clock together as well as attach arcs that connect carefully chosen vertices on these hands to obtain a *double clock*. The double clock is a gadget that both ensures that two clocks show the exact same time and that this time corresponds to the selection of two indices whose sum is *exactly*  $n + 1$ . Roughly speaking, it is mentally convenient to associate each double clock with a different *time zone* that encodes the selection of *one* element. Here, since our source problem is a graph problem, the natural choice of an element is a vertex. Having established a time zone for each selection of one element, we turn to *synchronize* hands of *different* double clocks. For this purpose, we introduce the *synchronizer*, which is a gadget that resembles a *folded grid*. We remark that this specific gadget is different yet inspired by a folded grid gadget that is the core of the paper [46]. Having double clocks and synchronizers at hand, we are finally able to present the entire reduction in an intuitive (yet

precise) manner. Lastly, we prove that our reduction is correct. At this point, having already established key properties of our gadgets, the reverse direction (“solution to DOCT  $\rightarrow$  solution to PSI”) is simple. For the forward direction (“solution to PSI  $\rightarrow$  solution to DOCT”) we exhibit a partition of the vertex set of the output digraph into pairwise-disjoint sets on which we can define a topological order, such that the graph induced by each set can be shown to exclude directed odd cycles.<sup>6</sup>

### 3.1 Annotations

Let us begin our proof by integrating arc and vertex annotations into the definition of DOCT. More precisely, we generalize DOCT as follows.

ANNOTATED DOCT (A-DOCT)	
<i>Input:</i>	A digraph $D$ , a non-negative integer $k$ , a labeling function $\ell : A(D) \rightarrow \{0, 1\}$ , and a weight function $w : V(D) \rightarrow [2k + 1]$ .
<i>Question:</i>	Does there exist a subset $X \subseteq V(D)$ such that $w(X) \leq k$ and $X$ intersects every directed cycle $C$ of $D$ where $\ell(E(C))$ is odd?

Henceforth, in the context of A-DOCT, the term *directed odd cycle* would refer to a directed cycle such that  $\ell(E(C))$  is odd. As we show in this section, it is easy to see that in order to prove Theorems 1 and 3, we can focus on the A-DOCT problem.

Let us now present our reduction from A-DOCT to DOCT. For this purpose, let  $(D, k, \ell, w)$  be an instance of A-DOCT. Then, we construct an instance  $\mathbf{red}(D, k, \ell, w) = (D', k')$  of DOCT as follows. First, set  $k' = k$ . Let  $A_0 = \{a \in A(D) : \ell(a) = 0\}$  and  $A_1 = \{a \in A(D) : \ell(a) = 1\}$ . Next, define  $V(D') = P \cup Q$ , where  $P = \{p_a^i : i \in [\alpha k + 1], a \in A_0\}$  and  $Q = \{q_v^i : i \in [w(v)], v \in V(D)\}$ . Finally, we define  $A(D') = S \cup T \cup R$ , where  $S = \{(q_v^i, p_a^j) : q_v^i \in Q, p_a^j \in P, v \text{ is the tail of } a\}$ ,  $T = \{(p_a^i, q_v^j) : p_a^i \in P, q_v^j \in Q, v \text{ is the head of } a\}$  and  $R = \{(q_u^i, q_v^j) : (u, v) \in A_1\}$ . Clearly,  $(D', k')$  can be computed in polynomial time. Notice that this reduction corresponds to first subdividing arcs in  $A_0$  once each, then replacing every original vertex with its weight-many copies, and finally replacing every new vertex with  $\alpha k + 1$  copies.

**Lemma 3.1.** *Let  $(D, k, \ell, w)$  be an instance of A-DOCT. If there exists a solution for  $(D, k, \ell, w)$ , then there exists a solution for  $\mathbf{red}(D, k, \ell, w) = (D', k')$ . Moreover, if there exists an  $\alpha$ -approximate solution for  $(D', k')$ , then there exists an  $\alpha$ -approximate solution for  $(D, k, \ell, w)$ .*

*Proof.* Fix  $\alpha \geq 1$ . In the first direction, let  $X$  be a solution for  $(D, k, \ell, w)$ . We claim that  $X' = \{q_v^i : v \in X, i \in [w(v)]\}$  is a solution to  $(D', k')$ . Suppose, by way of contradiction, that this claim is false. Then, since  $|X'| = w(X) \leq k = k'$ , there exists a directed odd cycle  $C'$  of minimum size of  $D' \setminus X'$ . If there exist  $q_v^i, q_v^j \in V(C') \cap Q$  such that  $i \neq j$ , then we obtain a contradiction to the choice of  $C'$ . Indeed, if we replace  $q_v^i$  by  $q_v^j$  in  $C'$ , then the result is a directed odd closed walk. Since a directed odd closed walk contains a directed odd cycle, we obtain a directed odd cycle that is shorter than  $C'$ . Hence, by the definitions of  $P$ ,  $S$  and  $T$ , we have that the graph  $C$  on  $\{v : q_v^i \in V(C')\}$ , where  $(u, v) \in A(C)$  if and only if there exist indices  $i, j, t$  such that either  $(q_u^i, q_v^j) \in A(C')$  or  $(q_u^i, p_{(q_u^i, q_v^j)}^t) \in A(C')$ , is a directed odd cycle of  $D \setminus X$ . Thus, we have reached a contradiction to the supposition that  $X$  is a solution to  $(D, k, \ell, w)$ .

Second, let  $X'$  be an  $\alpha$ -approximate solution for  $(D', k')$ . Without loss of generality, assume that  $X'$  is a minimal solution. We first claim that  $X' \cap P = \emptyset$ . For all  $a \in D(A)$ , the vertices  $p_a^i$  have the same set of outgoing neighbors and the same set of incoming neighbors. Hence, if there exist  $i \neq j$  such that  $p_a^i \in X'$  but  $p_a^j \notin X'$ , then  $X' \setminus \{p_a^i\}$  is also a solution to  $(D', k')$ .

<sup>6</sup>For the sake of clarity, we integrate the lemmata necessary to exhibit this partition into the sections presenting individual gadgets.

Indeed, if  $D' \setminus (X' \setminus \{p_a^i\})$  contains a directed odd cycle, then this cycle must contain  $p_a^i$ . By replacing  $p_a^i$  by  $p_a^j$ , we obtain a directed odd walk of  $D'$  (which contains a directed odd cycle of  $D'$ ). Hence, we reach a contradiction to the minimality of  $X'$ . Since there are  $\alpha k + 1$  vertices  $p_a^i$  while  $|X'| \leq \alpha k$ , we conclude that  $X' \cap P = \emptyset$ . Moreover, for all  $v \in V(D)$ , the vertices  $q_v^i$  have the same set of outgoing neighbors and the same set of incoming neighbors. Hence, we again deduce that there cannot exist  $i \neq j$  such that  $q_v^i \in X'$  but  $q_v^j \notin X'$ . Let us denote  $X = \{v : q_v^i \in X'\}$ . Then,  $w(X) = |X'| \leq \alpha k' = \alpha k$ . We claim that  $X$  is a solution to  $(D, k, \ell, w)$ . Suppose, by way of contradiction, that this claim is false. Then, let  $C$  be a directed odd cycle of  $D$ . Let  $C'$  be obtained from  $C$  by replacing each arc  $a = (u, v) \in A(C) \cap A_0$  by the two arcs  $(q_u^1, p_a^1)$  and  $(p_a^1, q_v^1)$ . By the definitions of  $P$ ,  $S$  and  $T$ , and since we have argued that  $X' \cap P = \emptyset$ , we have that  $C$  is a directed odd cycle of  $D' \setminus X'$ . Hence, we have reached a contradiction to the supposition that  $X'$  is an  $\alpha$ -approximate solution to  $(D', k')$ .  $\square$

As a corollary to Lemma 3.1, we derive the following result.

**Corollary 1.** *For all  $\alpha \geq 1$ , if there exists an  $\alpha$ -approximation algorithm for DOCT that runs in time  $\tau$ , then there exists an  $\alpha$ -approximation algorithm for A-DOCT that runs in time  $\mathcal{O}(\tau + n^{\mathcal{O}(1)})$ .*

*Proof.* If there exists an  $\alpha$ -approximation algorithm  $\mathcal{A}$  for DOCT that runs in time  $\tau$ , then we define  $\mathcal{B}$  as the algorithm that given an instance  $(D, k, \ell, w)$  of A-DOCT, constructs the instance  $\mathbf{red}(D, k, \ell, w) = (D', k')$  of DOCT, and calls algorithm  $\mathcal{A}$  with  $(D', k')$  as input. If  $\mathcal{A}$  returns an  $\alpha$ -approximate solution for  $(D', k')$ , we have shown (in Lemma 3.1) how to translate it to an  $\alpha$ -approximate solution for  $(D, k, \ell, w)$ . Moreover, if  $(D, k, \ell, w)$  is a yes-instance, then we have shown that  $(D', k')$  is a yes-instance. Hence,  $\mathcal{A}$  would return an  $\alpha$ -approximate solution for  $(D', k')$ . Thus, we obtain an  $\alpha$ -approximation algorithm for A-DOCT that runs in time  $\mathcal{O}(\tau + n^{\mathcal{O}(1)})$ .  $\square$

## 3.2 The Basic Clock Gadget

Let  $n, k \in \mathbb{N}$  such that  $k \geq 100$ . Here, we define an  $(n, k)$ -forward clock and an  $(n, k)$ -reverse clock. Since  $n$  and  $k$  would be clear from context, we simply write *forward clock* and *reverse clock* rather than  $(n, k)$ -forward clock and  $(n, k)$ -reverse clock, respectively.

### 3.2.1 Forward Clock

**Structure.** We first define a forward clock  $C$ . The *face* of  $C$  is an “undirected” cycle, in the sense that consecutive vertices have arcs in both directions. The vertex set of  $C$  is the union of four pairwise-disjoint sets  $\widehat{R}$  (red),  $\widehat{B}$  (blue),  $\widehat{T}$  (time) and  $\{x\}$ . We refer the reader to Fig. 1. We set  $\widehat{R} = \{\widehat{r}_i : i \in [n]_0\}$ ,  $\widehat{B} = \{\widehat{b}_i : i \in [n]_0\}$  and  $\widehat{T} = \{\widehat{t}_{i, n-i-1} : i \in [n-1]_0\}$ . The arc set of the face is the union of the following three pairwise-disjoint sets.

- $\{(x, \widehat{r}_n), (x, \widehat{b}_n), (\widehat{r}_n, x), (\widehat{b}_n, x)\}$ .
- $\{(\widehat{b}_i, \widehat{r}_{n-i}) : i \in [n]_0\} \cup \{(\widehat{r}_{n-i}, \widehat{b}_i) : i \in [n]_0\}$ .
- $\{(\widehat{r}_i, \widehat{t}_{i, n-i-1}) : i \in [n-1]_0\} \cup \{(\widehat{b}_{n-i-1}, \widehat{t}_{i, n-i-1}) : i \in [n-1]_0\} \cup \{(\widehat{t}_{i, n-i-1}, \widehat{r}_i) : i \in [n-1]_0\} \cup \{(\widehat{t}_{i, n-i-1}, \widehat{b}_{n-i-1}) : i \in [n-1]_0\}$ .

The *hands* of  $C$  are two directed paths, red and blue (see Fig. 2). The vertex set of the red path is the union of two pairwise disjoint sets,  $R = \{r_i : i \in [n]_0\}$  (red) and  $P = \{p_i : i \in [n]\}$  (pink). For all  $i \in [n]$ , we denote  $\text{pre}(p_i) = r_{i-1}$  and  $\text{post}(p_i) = r_i$ . The arc set of the red

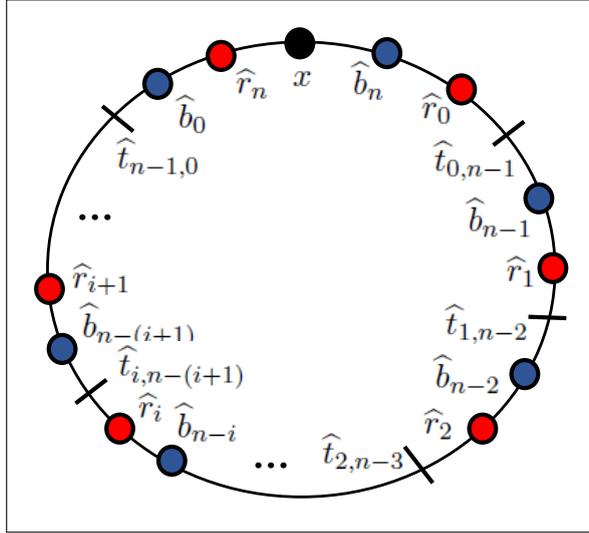


Figure 1: The face of a forward clock. The shapes used to represent nodes are chosen for the sake of clarity of the figures. Nodes that have the same shape serve a similar purpose. In particular, some vertices are represented by lines to emphasize that we would sometimes like to delete these vertices, which corresponds to “cutting” the cycle (or path) at this position. The colors distinguish between different types of nodes; red vertices are called  $r_i$ , blue vertices are called  $b_i$ , etc.

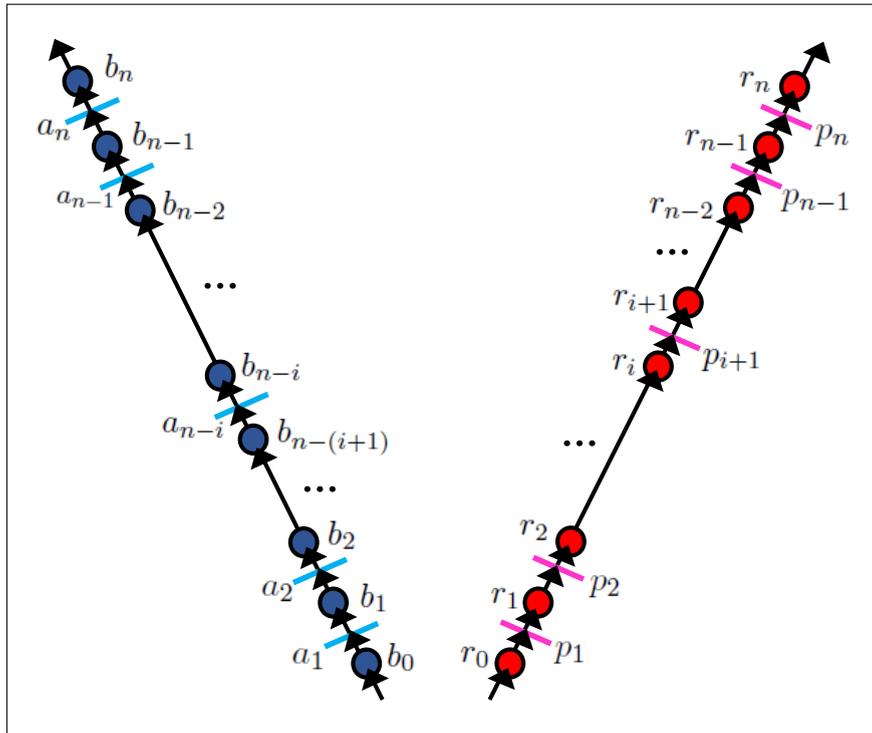


Figure 2: The hands of a forward clock together with the 4 arcs used to attach them. For all  $i \in [n]$ ,  $\text{pre}(p_i) = r_{i-1}$  and  $\text{post}(p_i) = r_i$ , and  $\text{pre}(a_i) = b_{i-1}$  and  $\text{post}(a_i) = b_i$ .

path is  $\{(\text{pre}(p_i), p_i) : i \in [n]\} \cup \{(p_i, \text{post}(p_i)) : i \in [n]\}$ . Symmetrically, the blue path is the union of two pairwise disjoint sets,  $B = \{b_i : i \in [n]_0\}$  (blue) and  $A = \{a_i : i \in [n]\}$  (azure). For all  $i \in [n]$ , we denote  $\text{pre}(a_i) = b_{i-1}$  and  $\text{post}(a_i) = b_i$ . The arc set of the blue path is  $\{(\text{pre}(a_i), a_i) : i \in [n]\} \cup \{(a_i, \text{post}(a_i)) : i \in [n]\}$ .

The hands are attached to the face as follows (see Fig. 3). First, we add the arcs  $(x, r_0)$  and  $(x, b_0)$ . Second, for all  $i \in [n]_0$ , we add the arcs  $(r_i, \widehat{r}_i)$  and  $(b_i, \widehat{b}_i)$ . Then, we “glue” the hands by adding a new vertex,  $y$ , and the arcs  $(r_n, y)$ ,  $(b_n, y)$  and  $(y, x)$ .

Finally, let us annotate  $C$  (see Fig. 3). The labels of the arcs in the set  $\{(x, \widehat{r}_n), (\widehat{r}_n, x), (y, x)\} \cup \{(b_i, \widehat{b}_i) : i \in [n]_0\}$  are equal to 1, and the labels of all other arcs are equal to 0. Moreover, the weight of the vertices in the set  $\widehat{T} \cup P \cup A$  are equal to 10, and the weights of all other vertices are equal to  $2k + 1$ . This completes the description of  $C$ . When the clock  $C$  is not clear from context, we add the notation  $(C)$  to an element (vertex set or vertex) of the clock. For example, we may write  $R(C)$  and  $x(C)$ .

**Intuition.** It is easy to verify that any solution has to have non-empty intersection with the face of the clock, the cycle starting at  $x$  following the blue hand to  $y$  and returning to  $x$ , and the cycle starting at  $x$  following the red hand to  $y$  and returning to  $x$ . This requires total weight at least 30. The following Lemmata establish the structure of all directed odd cycle transversals of the clock gadget of weight exactly 30, and prove that any directed odd cycle transversal of weight more than 30 must have weight at least 40.

A directed odd cycle transversal of the clock gadget of weight less than 40 must pick one vertex  $\widehat{t}_{s, n-(s+1)}$  on the face of the clock, one vertex  $p_i$  on the red hand, and one vertex  $a_j$  on the blue hand. We will show that these three vertices form a directed odd cycle transversal of the clock gadget if and only if  $i - 1 \leq s$  and  $j \leq n - s$ . In other words we can think of the variable  $s$  as the “time” shown by the clock, and as the time  $s$  increases we are allowed to pick the vertex  $p_i$  further away from  $x$  on the red hand, but at the same time we need to pick the vertex  $a_j$  closer to  $x$ . This is captured in Definition 3.1 and Lemmata 3.2 and 3.3.

**Properties.** By the definition of a forward clock, we directly identify which directed odd cycles are present in such a clock.

**Observation 3.1.** *Let  $C$  be a forward clock. The set of directed odd cycles of  $C$  is the union of the following sets.*

- **Type 1:** *The set whose only directed odd cycle is the one consisting of the entire red hand and the arc  $(y, x)$ .*
- **Type 2:** *The set whose only directed odd cycle is the one consisting of the entire blue hand and the arc  $(y, x)$ .*
- **Type 3:** *For all  $i \in [n]_0$ , this set contains the directed odd cycle consisting of the directed path from  $x$  to  $r_i$  on the red hand, the arc  $(r_i, \widehat{r}_i)$ , and the directed path from  $\widehat{r}_i$  to  $x$  on the face of the clock that contains the arc  $(\widehat{r}_n, x)$ .*
- **Type 4:** *For all  $i \in [n]_0$ , this set contains the directed odd cycle consisting of the directed path from  $x$  to  $b_i$  on the blue hand, the arc  $(b_i, \widehat{b}_i)$ , and the directed path from  $\widehat{b}_i$  to  $x$  on the face of the clock that contains the arc  $(\widehat{b}_n, x)$ .*
- **Type 5:** *The set whose only directed odd cycle is the face of the clock.*

We proceed to derive properties of “cuts” of a forward clock. To this end, we first need to define the kind of sets using which we would like to “cut” forward clocks.

**Definition 3.1.** *Let  $C$  be a forward clock. We say that a set  $X \subseteq V(C)$  cuts  $C$  precisely if there exist  $i, j, s \in [n]$  such that  $X = \{p_i, a_j, \widehat{t}_{s, n-s-1}\}$ ,  $i - 1 \leq s$  and  $j \leq n - s$ .*

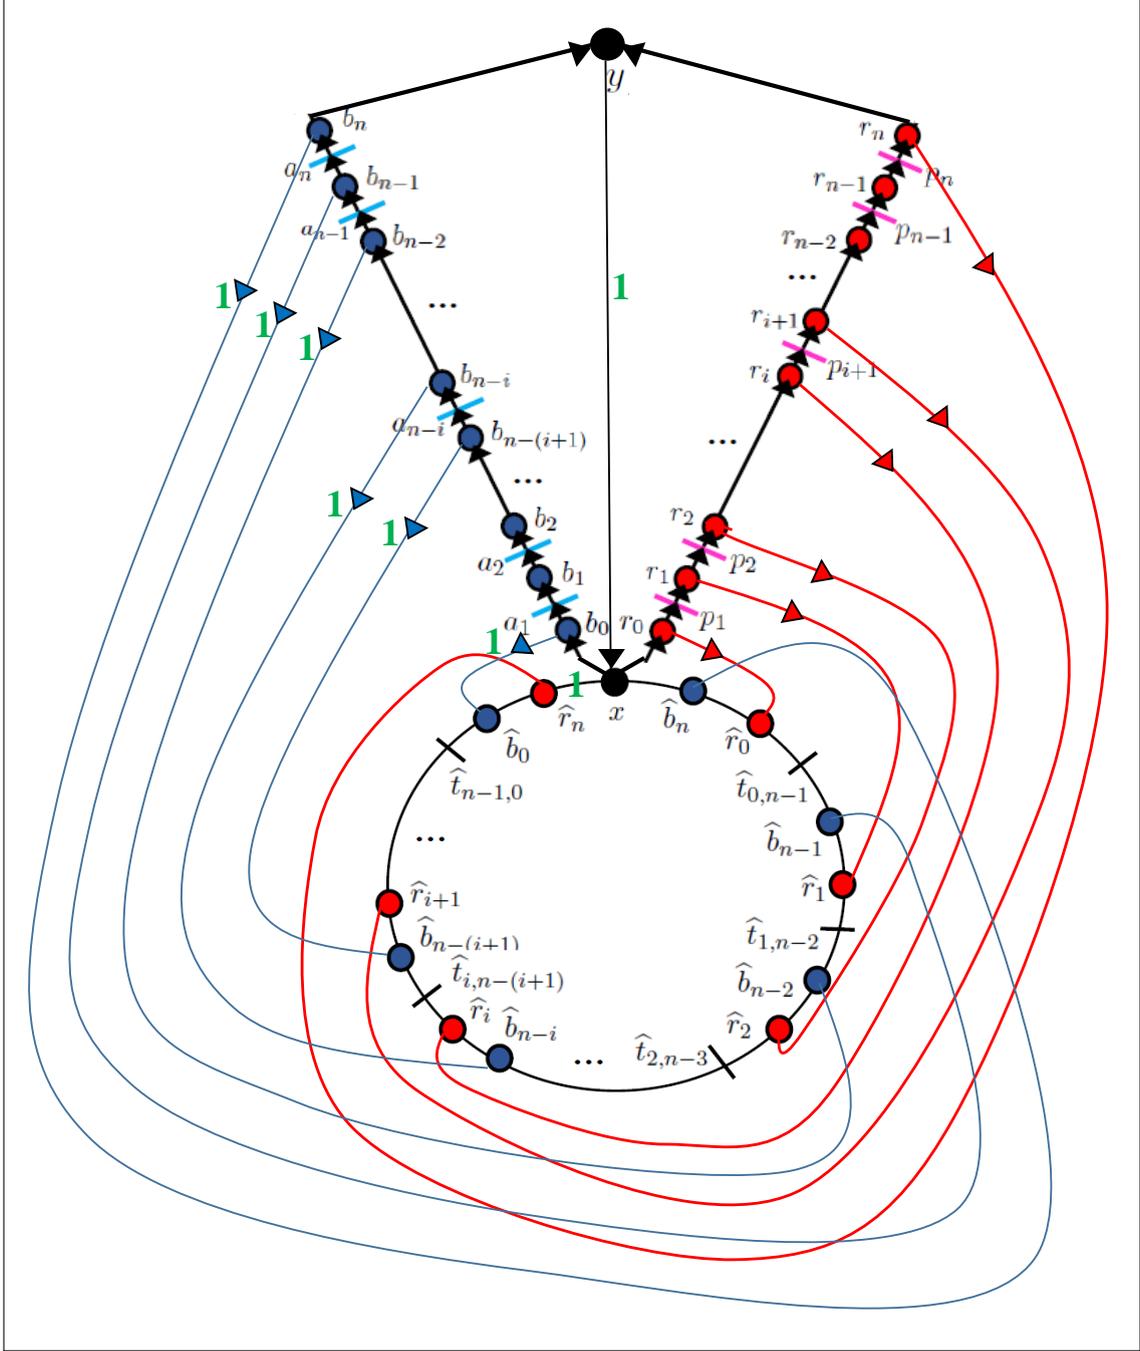


Figure 3: A forward clock. The arcs labeled 1 are marked by a green ‘1’. The weight of vertices marked by circles is  $2k + 1$ , and the weight of vertices marked by lines is 10.

Definition 3.1 directly implies the following observation.

**Observation 3.2.** *Let  $C$  be a forward clock. If  $X = \{p_i, a_j, \hat{t}_{s,n-s-1}\}$  is a set that cuts  $C$  precisely, then  $i + j \leq n + 1$ .*

We are now ready to present the desired properties of “cuts” of a forward clock.

**Lemma 3.2.** *Let  $C$  be a forward clock. A set  $X \subseteq V(C)$  is a directed odd cycle transversal of  $C$  of weight exactly 30 if and only if  $X$  cuts  $C$  precisely.*

*Proof.* In the forward direction, let  $X \subseteq V(C)$  be a directed odd cycle transversal of  $C$  of weight exactly 30. Recall that  $2k + 1 > 40$ . Hence, to intersect the directed odd cycle of Type 1 (see Observation 3.1), the set  $X$  must contain a vertex of the form  $p_i$  for some  $i \in [n]$ . Symmetrically, to intersect the directed odd cycle of Type 2, the set  $X$  must contain a vertex of the form  $a_j$  for some  $j \in [n]$ . Moreover, to intersect the directed odd cycle of Type 5, the set  $X$  must contain a vertex of the form  $\hat{t}_{s,n-s-1}$  for some  $s \in [n-1]_0$ . Since  $w(X) = 30$ , we deduce that  $X = \{p_i, a_j, \hat{t}_{s,n-s-1}\}$ . To prove that  $X$  cuts  $C$  precisely, it remains to show that  $i - 1 \leq s$  and  $j \leq n - s$ . For this purpose, consider the directed odd cycle of Type 3 that consists of the directed path from  $x$  to  $r_{i-1}$  on the red hand, the arc  $(r_{i-1}, \hat{r}_{i-1})$ , and the directed path from  $\hat{r}_{i-1}$  to  $x$  on the face of the clock that contains the arc  $(\hat{r}_n, x)$ . Since  $X$  must intersect this directed odd cycle, it must hold that  $i - 1 \leq s$ . Now, consider the directed odd cycle of Type 4 that consists of the directed path from  $x$  to  $b_{j-1}$  on the blue hand, the arc  $(b_{j-1}, \hat{b}_{j-1})$ , and the directed path from  $\hat{b}_{j-1}$  to  $x$  on the face of the clock that contains the arc  $(\hat{b}_n, x)$ . Since  $X$  must intersect this directed odd cycle, it must also hold that  $j - 1 \leq n - s - 1$ , and therefore  $j \leq n - s$ .

In the reverse direction, let  $X \subseteq V(C)$  be a set that cuts  $C$  precisely. Then, there exist  $i, j, s \in [n]$  such that  $X = \{p_i, a_j, \hat{t}_{s,n-s-1}\}$ ,  $i - 1 \leq s$  and  $j \leq n - s$ . Clearly,  $w(X) = 30$ . Since  $p_i \in X$ , it holds that  $X$  intersects the directed odd cycle of Type 1 as well as every directed odd cycle of Type 3 that consists of a directed path from  $x$  to  $r_{i'}$  on the red hand for some  $i' \geq i$ , the arc  $(r_{i'}, \hat{r}_{i'})$ , and the directed path from  $\hat{r}_{i'}$  to  $x$  on the face of the clock that contains the arc  $(\hat{r}_n, x)$ . Symmetrically, since  $a_j \in X$ , it holds that  $X$  intersects the directed odd cycle of Type 2 as well as every directed odd cycle of Type 4 that consists of a directed path from  $x$  to  $b_{j'}$  on the blue hand for some  $j' \geq j$ , the arc  $(b_{j'}, \hat{b}_{j'})$ , and the directed path from  $\hat{b}_{j'}$  to  $x$  on the face of the clock that contains the arc  $(\hat{b}_n, x)$ . Since  $\hat{t}_{s,n-s-1} \in X$ , it holds that  $X$  intersects that directed odd cycle of Type 5. Moreover, since  $i - 1 \leq s$ , it holds that  $X$  intersects every directed odd cycle of Type 3 that consists of a directed path from  $x$  to  $r_{i'}$  on the red hand for some  $i' < i$ , the arc  $(r_{i'}, \hat{r}_{i'})$ , and the directed path from  $\hat{r}_{i'}$  to  $x$  on the face of the clock that contains the arc  $(\hat{r}_n, x)$ . Symmetrically, since  $j - 1 \leq n - s - 1$ , it holds that  $X$  intersects every directed odd cycle of Type 4 that consists of a directed path from  $x$  to  $b_{j'}$  on the blue hand for some  $j' < j$ , the arc  $(b_{j'}, \hat{b}_{j'})$ , and the directed path from  $\hat{b}_{j'}$  to  $x$  on the face of the clock that contains the arc  $(\hat{b}_n, x)$ . We have thus verified that  $X$  is a directed odd cycle transversal of  $C$ .  $\square$

**Lemma 3.3.** *Let  $C$  be a forward clock. The weight of a set  $X \subseteq V(C)$  that is a directed odd cycle transversal of  $C$  but does not cut  $C$  precisely is at least 40.*

*Proof.* Consider the directed odd cycles of  $C$  of Types 1, 2 and 5 (see Observation 3.1). The only vertices that are present in at least two of these cycles are of weight  $2k + 1 > 40$ . Hence, if  $w(X) \leq 40$ , then  $X$  must contain at least three vertices of  $C$ , each of weight 10. The set  $X$  cannot contain exactly three vertices of  $C$  of weight 10, since then  $w(X) = 30$ , in which Lemma 3.2 implies  $X$  should have cut  $C$  precisely. Hence, if  $w(X) \leq 40$ , then  $X$  contains at least four vertices of  $C$ , and therefore  $w(X) \geq 40$ .  $\square$

### 3.2.2 Reverse Clock

A reverse clock is simply a forward clock where the directions of *all* arcs have been reversed. For readability reasons it is useful to re-name the vertices (so that, for example, we avoid arcs going from *post* to *pre*), and modify the numbering of the vertices. The proofs in this subsection are identical to the proofs in Subsection 3.2.1, up to permutation of vertex names, changing

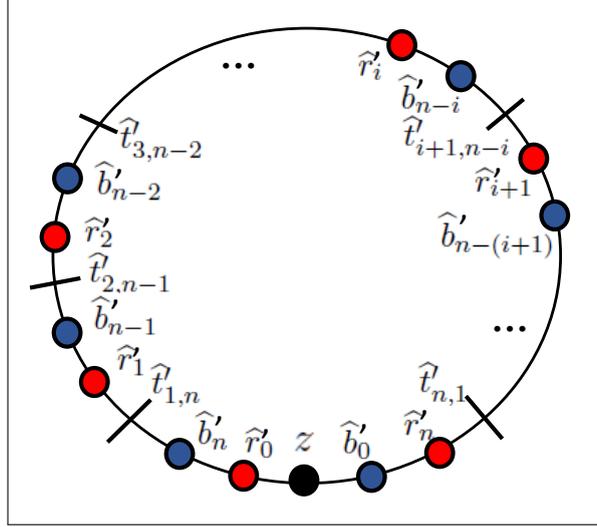


Figure 4: The face of a reverse clock.

“from” to “to” and vice versa, and changing the indices. We include the proofs and figures for completeness and ease of reference, since changing the indices slightly changes the equations.

**Structure.** The *face* of a reverse clock  $C$  is an “undirected” cycle whose vertex set is the union of four pairwise-disjoint sets  $\hat{R}'$  (red),  $\hat{B}'$  (blue),  $\hat{T}'$  (time) and  $\{z\}$ . We refer the reader to Fig. 4. We set  $\hat{R}' = \{\hat{r}'_i : i \in [n]_0\}$ ,  $\hat{B}' = \{\hat{b}'_i : i \in [n]_0\}$  and  $\hat{T}' = \{\hat{t}'_{i,n-i+1} : i \in [n]\}$ . The arc set of the face is the union of the following three pairwise-disjoint sets.

- $\{(z, \hat{r}'_0), (z, \hat{b}'_0), (\hat{r}'_0, z), (\hat{b}'_0, z)\}$ .
- $\{(\hat{b}'_i, \hat{r}'_{n-i}) : i \in [n]_0\} \cup \{(\hat{r}'_{n-i}, \hat{b}'_i) : i \in [n]_0\}$ .
- $\{(\hat{r}'_i, \hat{t}'_{i,n-i+1}) : i \in [n]\} \cup \{(\hat{b}'_{n-i+1}, \hat{t}'_{i,n-i+1}) : i \in [n]\} \cup \{(\hat{t}'_{i,n-i+1}, \hat{r}'_i) : i \in [n]\} \cup \{(\hat{t}'_{i,n-i+1}, \hat{b}'_{n-i+1}) : i \in [n]\}$ .

The *hands* of  $C$  are two directed paths, red and blue. These hands are defined exactly as the hands of a forward clock, except that tags are added to the names of all of their vertices (see Fig. 5). The hands are attached to the face as follows (see Fig. 6). First, we add the arcs  $(\hat{r}'_n, z)$  and  $(\hat{b}'_n, z)$ . Second, for all  $i \in [n]_0$ , we add the arcs  $(\hat{r}'_i, \hat{r}'_i)$  and  $(\hat{b}'_i, \hat{b}'_i)$ . Then, we “glue” the hands by adding a new vertex,  $y$ , and the arcs  $(y, \hat{r}'_0)$ ,  $(y, \hat{b}'_0)$  and  $(z, y)$ .

Finally, let us annotate  $C$  (see Fig. 6). The labels of the arcs in the set  $\{(z, \hat{r}'_0), (\hat{r}'_0, z), (z, y)\} \cup \{(\hat{b}'_i, \hat{b}'_i) : i \in [n]_0\}$  are equal to 1, and the labels of all other arcs are equal to 0. Moreover, the weight of the vertices in the set  $\hat{T}' \cup P' \cup A'$  are equal to 10, and the weights of all other vertices are equal to  $2k + 1$ . This completes the description of  $C$ .

**Properties.** By the definition of a reverse clock, we directly identify which directed odd cycles are present in such a clock.

**Observation 3.3.** *Let  $C$  be a reverse clock. The set of directed odd cycles of  $C$  is the union of the following sets.*

- **Type 1:** *The set whose only directed odd cycle is the one consisting of the entire red hand and the arc  $(z, y)$ .*
- **Type 2:** *The set whose only directed odd cycle is the one consisting of the entire blue hand and the arc  $(z, y)$ .*

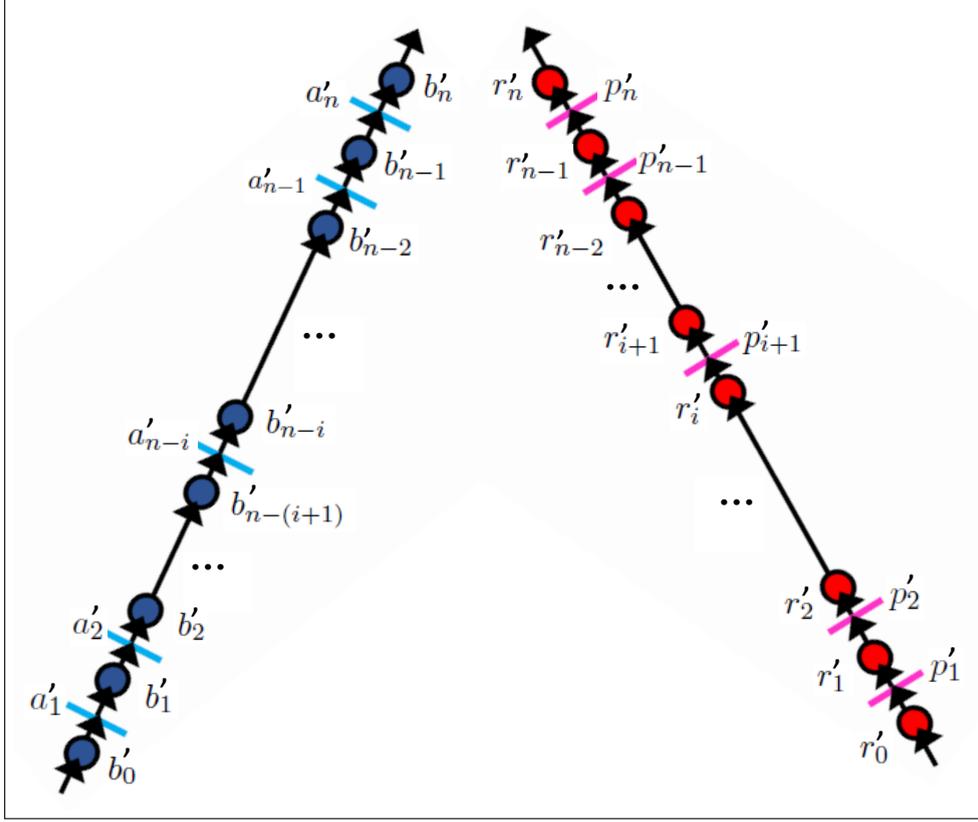


Figure 5: The hands of a reverse clock. For all  $i \in [n]$ ,  $\text{pre}(p'_i) = r'_{i-1}$  and  $\text{post}(p'_i) = r'_i$ , and  $\text{pre}(a'_i) = b'_{i-1}$  and  $\text{post}(a'_i) = b'_i$ .

- **Type 3:** For all  $i \in [n]_0$ , this set contains the directed odd cycle consisting of the arc  $(\widehat{r}'_i, r'_i)$ , the directed path from  $r'_i$  to  $z$  on the red hand, and the directed path from  $z$  to  $\widehat{r}'_i$  on the face of the clock that contains the arc  $(z, \widehat{r}'_0)$ .
- **Type 4:** For all  $i \in [n]_0$ , this set contains the directed odd cycle consisting of the arc  $(\widehat{b}'_i, b'_i)$ , the directed path from  $b'_i$  to  $z$  on the blue hand, and the directed path from  $z$  to  $\widehat{b}'_i$  on the face of the clock that contains the arc  $(z, \widehat{b}'_0)$ .
- **Type 5:** The set whose only directed odd cycle is the face of the clock.

As in the case of a forward clock, we proceed to derive properties of “cuts” of a reverse clock. To this end, we first need to define the kind of sets using which we would like to “cut” reverse clocks.

**Definition 3.2.** Let  $C$  be a reverse clock. We say that a set  $X \subseteq V(C)$  cuts  $C$  precisely if there exist  $i, j, s \in [n]$  such that  $X = \{p'_i, a'_j, \widehat{t}_{s, n-s+1}\}$ ,  $s \leq i$  and  $n - s + 1 \leq j$ .

Definition 3.2 directly implies the following observation. Note that due to our placement of indices, the inequality is complementary to the one in Observation 3.2.

**Observation 3.4.** Let  $C$  be a reverse clock. If  $X = \{p'_i, a'_j, \widehat{t}_{s, n-s-1}\}$  is a set that cuts  $C$  precisely, then  $i + j \geq n + 1$ .

We are now ready to present the desired properties of “cuts” of a reverse clock.

**Lemma 3.4.** Let  $C$  be a reverse clock. A set  $X \subseteq V(C)$  is a directed odd cycle transversal of  $C$  of weight exactly 30 if and only if  $X$  cuts  $C$  precisely.

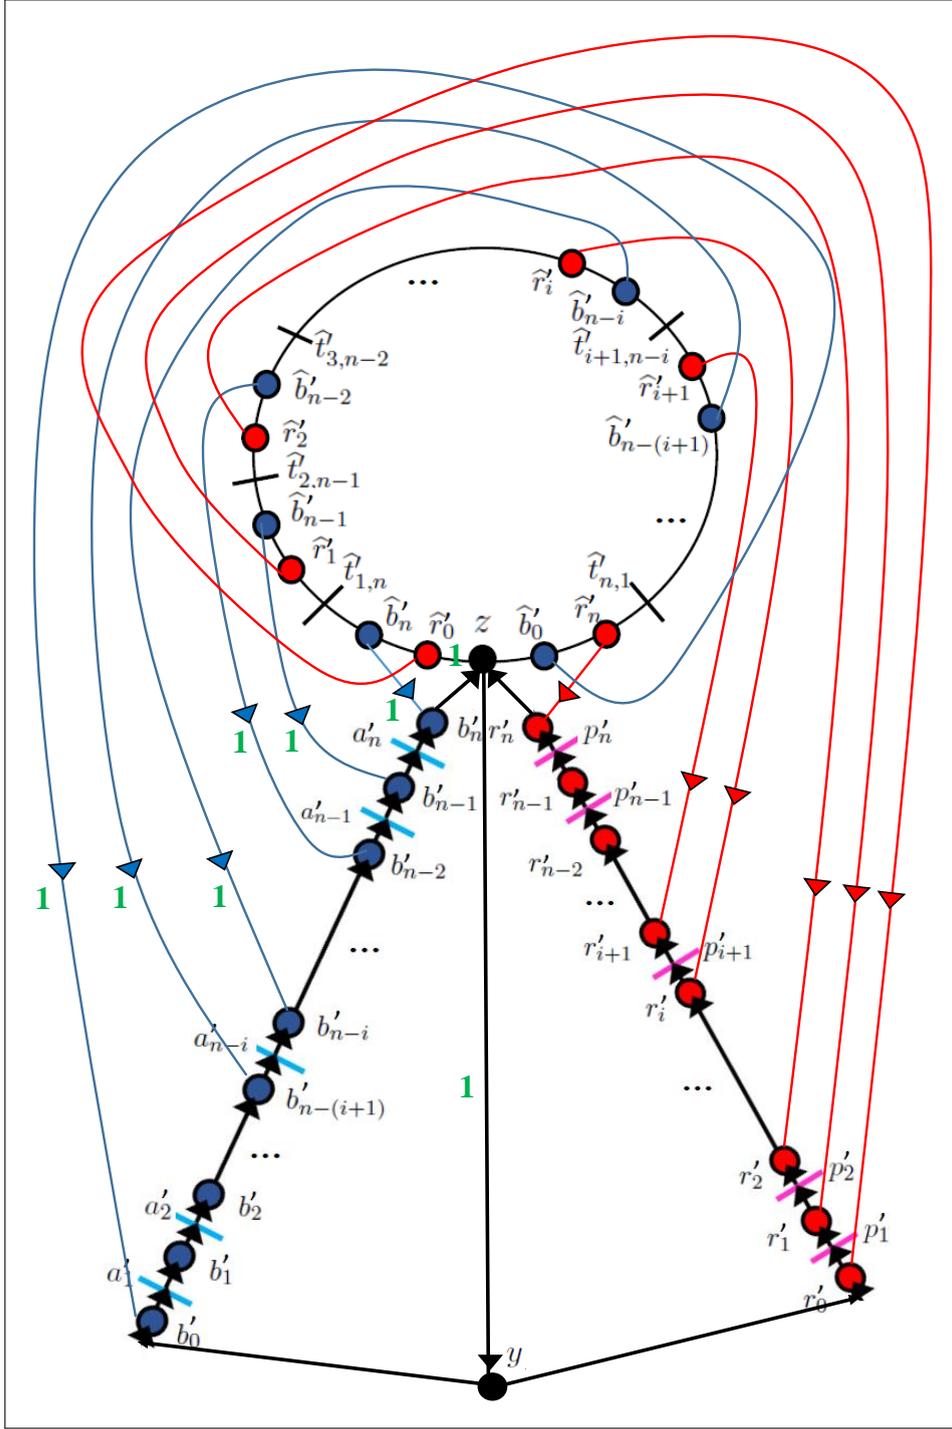


Figure 6: A reverse clock. The arcs labeled 1 are marked by a green ‘1’. The weight of vertices marked by circles is  $2k + 1$ , and the weight of vertices marked by lines is 10.

*Proof.* In the forward direction, let  $X \subseteq V(C)$  be a directed odd cycle transversal of  $C$  of weight exactly 30. Recall that  $2k + 1 > 40$ . Hence, to intersect the directed odd cycle of Type 1 (see Observation 3.3), the set  $X$  must contain a vertex of the form  $p'_i$  for some  $i \in [n]$ . Symmetrically, to intersect the directed odd cycle of Type 2, the set  $X$  must contain a vertex of the form  $a'_j$  for some  $j \in [n]$ . Moreover, to intersect the directed odd cycle of Type 5, the set  $X$  must contain a vertex of the form  $\hat{t}'_{s,n-s+1}$  for some  $s \in [n]$ . Since  $w(X) = 30$ , we deduce

that  $X = \{p'_i, a'_j, \widehat{t}_{s, n-s+1}\}$ . To prove that  $X$  cuts  $C$  precisely, it remains to show that  $s \leq i$  and  $n - s + 1 \leq j$ . For this purpose, consider the directed odd cycle of Type 3 that consists of the arc  $(\widehat{r}'_i, r'_i)$ , the directed path from  $r'_i$  to  $z$  on the red hand, and the directed path from  $z$  to  $\widehat{r}'_i$  on the face of the clock that contains the arc  $(z, \widehat{r}'_0)$ . Since  $X$  must intersect this directed odd cycle, it must hold that  $s \leq i$ . Now, consider the directed odd cycle of Type 4 that consists of the arc  $(\widehat{b}'_j, b'_j)$ , the directed path from  $b'_j$  to  $z$  on the blue hand, and the directed path from  $z$  to  $\widehat{b}'_j$  on the face of the clock that contains the arc  $(z, \widehat{b}'_0)$ . Since  $X$  must intersect this directed odd cycle, it must also hold that  $n - s + 1 \leq j$ .

In the reverse direction, let  $X \subseteq V(C)$  be a set that cuts  $C$  precisely. Then, there exist  $i, j, s \in [n]$  such that  $X = \{p'_i, a'_j, \widehat{t}_{s, n-s+1}\}$ ,  $s \leq i$  and  $n - s + 1 \leq j$ . Clearly,  $w(X) = 30$ . Since  $p'_i \in X$ , it holds that  $X$  intersects the directed odd cycle of Type 1 as well as every directed odd cycle of Type 3 that consists of the arc  $(\widehat{r}'_{i'}, r'_{i'})$  for some  $i' < i$ , the directed path from  $r'_{i'}$  to  $z$  on the red hand, and the directed path from  $z$  to  $\widehat{r}'_{i'}$  on the face of the clock that contains the arc  $(z, \widehat{r}'_0)$ . Symmetrically, since  $a'_j \in X$ , it holds that  $X$  intersects the directed odd cycle of Type 2 as well as every directed odd cycle of Type 4 that consists of the arc  $(\widehat{b}'_{j'}, b'_{j'})$  for some  $j' < j$ , the directed path from  $b'_{j'}$  to  $z$  on the blue hand, and the directed path from  $z$  to  $\widehat{b}'_{j'}$  on the face of the clock that contains the arc  $(z, \widehat{b}'_0)$ . Since  $\widehat{t}_{s, n-s+1} \in X$ , it holds that  $X$  intersects that directed odd cycle of Type 5. Moreover, since  $s \leq i$ , it holds that  $X$  intersects every directed odd cycle of Type 3 that consists of the arc  $(\widehat{r}'_{i'}, r'_{i'})$  for some  $i' \geq i$ , the directed path from  $r'_{i'}$  to  $z$  on the red hand, and the directed path from  $z$  to  $\widehat{r}'_{i'}$  on the face of the clock that contains the arc  $(z, \widehat{r}'_0)$ . Symmetrically, since  $n - s + 1 \leq j$ , it holds that  $X$  intersects every directed odd cycle of Type 4 that consists of the arc  $(\widehat{b}'_{j'}, b'_{j'})$  for some  $j' \geq j$ , the directed path from  $b'_{j'}$  to  $z$  on the blue hand, and the directed path from  $z$  to  $\widehat{b}'_{j'}$  on the face of the clock that contains the arc  $(z, \widehat{b}'_0)$ . We have thus verified that  $X$  is a directed odd cycle transversal of  $C$ .  $\square$

**Lemma 3.5.** *Let  $C$  be a reverse clock. The weight of a set  $X \subseteq V(C)$  that is a directed odd cycle transversal of  $C$  but does not cut  $C$  precisely is at least 40.*

*Proof.* Consider the directed odd cycles of  $C$  of Types 1, 2 and 5 (see Observation 3.3). The only vertices that are present in at least two of these cycles are of weight  $2k + 1 > 40$ . Hence, if  $w(X) \leq 40$ , then  $X$  must contain at least three vertices of  $C$ , each of weight 10. The set  $X$  cannot contain exactly three vertices of weight 10 of  $C$ , since then  $w(X) = 30$ , in which Lemma 3.4 implies  $X$  should have cut  $C$  precisely. Hence, if  $w(X) \leq 40$ , then  $X$  contains at least four vertices of  $C$ , and therefore  $w(X) \geq 40$ .  $\square$

### 3.3 The Double Clock Gadget

**Structure.** Roughly speaking, an  $(n, k)$ -double clock is the result of gluing the tips of the hands of an  $(n, k)$ -forward clock and an  $(n, k)$ -reverse clock together as well as adding a 1-labeled arc from every vertex on the hands of the reverse clock to its “twin” on the forward clock. In what follows, since  $n$  and  $k$  would be clear from context, we omit explicit references to  $(n, k)$ . Formally, a double clock  $\widetilde{C}$  is defined as the digraph obtained as follows. Let  $C$  be a forward clock, and let  $C'$  be a reverse clock. Identify the vertex  $y$  of both of these clocks (see Fig. 7). All other vertices are distinct. Now, for all  $i \in [n]_0$ , add the arcs  $(r'_i, r_i)$  and  $(b'_i, b_i)$ , and let the labels of both of these arcs be 1 (see Fig. 8).

**Properties.** By the definition of a double clock, we first directly identify which directed odd cycles are present in such a clock.

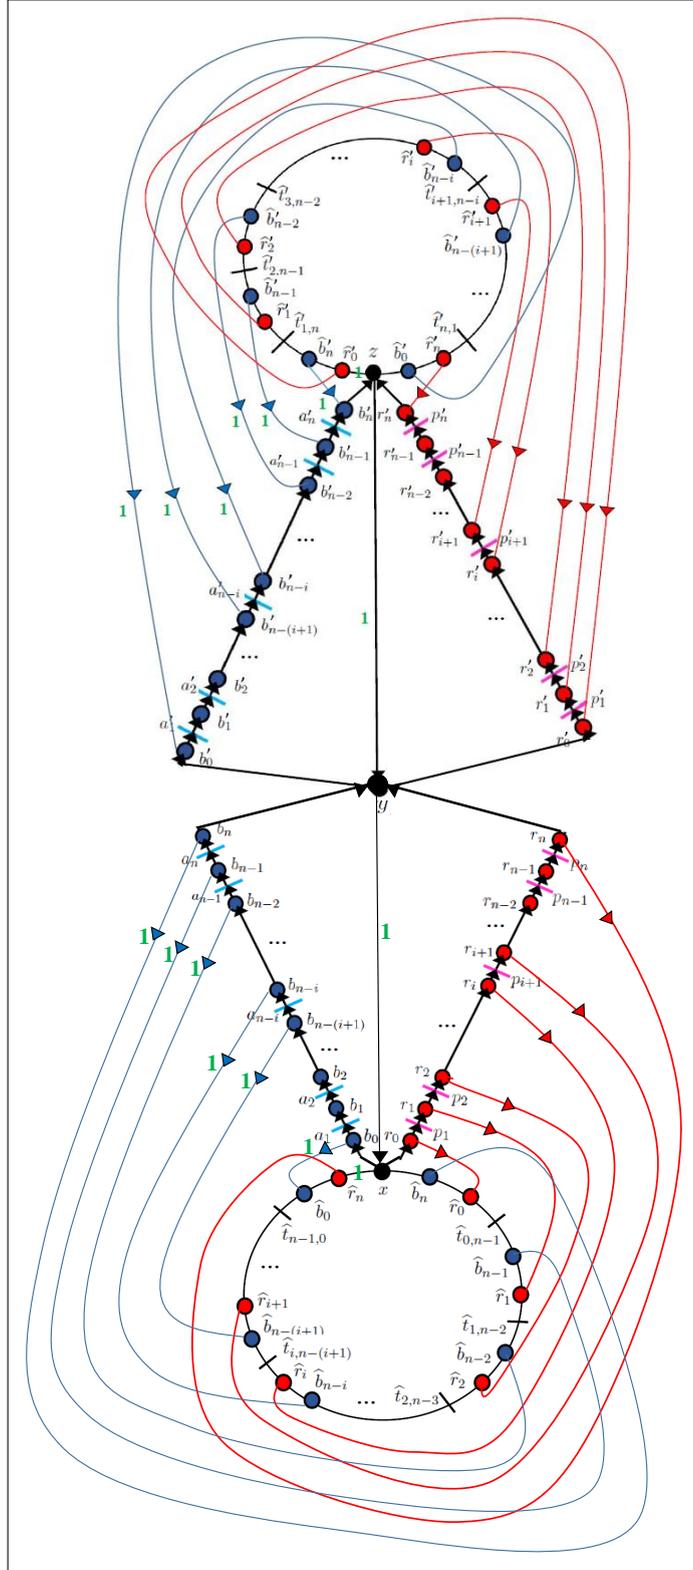


Figure 7: Gluing a forward clock and a reverse clock.

**Observation 3.5.** Let  $\tilde{C}$  be a double clock. The set of directed odd cycles of  $\tilde{C}$  is the union of the following sets.

- **Forward:** The set of directed odd cycles completely contained in the forward clock (see



- **Double Red:** For all  $i \in [n]_0$ , this set contains the direct odd cycle consisting of the arc  $(r'_i, r_i)$ , the directed path from  $r_i$  to  $y$  on the red hand of the forward clock, and the directed path from  $y$  to  $r'_i$  on the red hand of the reverse clock.
- **Double Blue:** For all  $i \in [n]_0$ , this set contains the direct odd cycle consisting of the arc  $(b'_i, b_i)$ , the directed path from  $b_i$  to  $y$  on the blue hand of the forward clock, and the directed path from  $y$  to  $b'_i$  on the blue hand of the reverse clock.

We proceed to derive properties of “cuts” of a double clock. To this end, we again first need to define the kind of sets using which we would like to “cut” double clocks.

**Definition 3.3.** Let  $\tilde{C}$  be a double clock. We say that a set  $X \subseteq V(\tilde{C})$  cuts  $\tilde{C}$  precisely if there exists  $i \in [n]$  such that  $X = \{p_i, a_{n-i+1}, p'_i, a'_{n-i+1}, \hat{t}_{i-1, n-i}, \hat{t}'_{i, n-i+1}\}$ .

We are now ready to present desired properties of “cuts” of a double clock.

**Lemma 3.6.** Let  $\tilde{C}$  be a double clock. A set  $X \subseteq V(\tilde{C})$  is a directed odd cycle transversal of  $\tilde{C}$  of weight exactly 60 if and only if  $X$  cuts  $\tilde{C}$  precisely.

*Proof.* In the forward direction, let  $X \subseteq V(\tilde{C})$  be a directed odd cycle transversal of  $\tilde{C}$  of weight exactly 60. By Lemmata 3.2, 3.3, 3.4 and 3.5, and since  $w(X) = 60$  while the only vertex that the forward and reverse clocks of  $\tilde{C}$  have in common is  $y$ , we deduce that  $X$  is of the following form. The set  $X$  is the union of two pairwise disjoint sets,  $Y$  and  $Y'$ , such that  $Y$  cuts the forward clock of  $\tilde{C}$  precisely, and  $Y'$  cuts the reverse clock of  $\tilde{C}$  precisely. Hence, there exist  $i, j, s \in [n]$  such that  $Y = \{p_i, a_j, \hat{t}_{s, n-s-1}\}$ ,  $i-1 \leq s$  and  $j \leq n-s$ . In particular, by Observation 3.2, it holds that  $i+j \leq n+1$ . Moreover, there exist  $i', j', s' \in [n]$  such that  $Y' = \{p'_{i'}, a'_{j'}, \hat{t}'_{s', n-s'+1}\}$ ,  $s' \leq i'$  and  $n-s'+1 \leq j'$ . In particular, by Observation 3.4, it holds that  $i'+j' \geq n+1$ . Thus, to prove the forward direction, it remains to show that  $i = i'$ ,  $j = j'$  and  $j = n-i+1$ . Indeed, if these equalities hold, then necessarily  $s = i-1$  and  $s' = i$ .

We claim that  $i \geq i'$ . Indeed, if  $i < i'$ , then  $X$  does not intersect the directed odd cycle of Type “Double Red” (see Observation 3.5) that consists of the arc  $(\text{post}(p'_i), \text{post}(p_i)) = (r'_i, r_i)$ , the directed path from  $r_i$  to  $y$  on the red hand of the forward clock, and the directed path from  $y$  to  $r'_i$  on the red hand of the reverse clock. Symmetrically, we claim that  $j \geq j'$ . Indeed, if  $j < j'$ , then  $X$  does not intersect the directed odd cycle of Type “Double Blue” that consists of the arc  $(\text{post}(a'_j), \text{post}(a_j)) = (b'_j, b_j)$ , the directed path from  $b_j$  to  $y$  on the blue hand of the forward clock, and the directed path from  $y$  to  $b'_j$  on the blue hand of the reverse clock.

Next, we observe that since  $i \geq i'$ ,  $j \geq j'$  and  $i'+j' \geq n+1$ , we have that  $i+j \geq n+1$ . However, since we have also argued that  $i+j \leq n+1$ , we have that  $i+j = n+1$ . Thus,  $j = n-i+1$ . Similarly, since  $i \geq i'$ ,  $j \geq j'$  and  $i+j \leq n+1$ , we have that  $i'+j' \leq n+1$ . However, since we have also argued that  $i'+j' \geq n+1$ , we have that  $i'+j' = n+1$ . Hence, we further deduce that  $i+j = i'+j'$ . However, since  $i \geq i'$  and  $j \geq j'$ , this implies that  $i = i'$  and  $j = j'$ . We thus conclude the correctness of the forward direction.

In the reverse direction, let  $X \subseteq V(\tilde{C})$  be a set that cuts  $\tilde{C}$  precisely. Then, there exist  $i \in [n]$  such that  $X = \{p_i, a_{n-i+1}, p'_i, a'_{n-i+1}, \hat{t}_{i-1, n-i}, \hat{t}'_{i, n-i+1}\}$ . Denote  $Y = \{p_i, a_{n-i+1}, \hat{t}_{i-1, n-i}\}$  and  $Y' = \{p'_i, a'_{n-i+1}, \hat{t}'_{i, n-i+1}\}$ . Observe that  $Y$  cuts the forward clock of  $\tilde{C}$  precisely, while  $Y'$  cuts the reverse clock of  $\tilde{C}$  precisely. Hence, by Lemmata 3.2 and 3.4, the set  $X$  intersects all directed odd cycles of Types “Forward” and “Reverse”. Thus, by Observation 3.5, and since it is clear that  $w(X) = 60$ , it remains to verify that  $X$  intersects all directed odd cycles of Types “Double Red” and “Double Blue”. Since  $p_i \in X$ , it holds that  $X$  intersects every directed odd cycle of Type “Double Red” that consists of the arc  $(r'_{i'}, r_{i'})$  for some  $i' < i$ , the directed path from  $r_{i'}$  to  $y$  on the red hand of the forward clock, and the directed path from  $y$  to  $r'_{i'}$  on the red hand of the reverse clock. Moreover, since  $p'_i \in X$ , it holds that  $X$  intersects every directed

odd cycle of Type “Double Red” that consists of the arc  $(r'_{i'}, r_{i'})$  for some  $i' \geq i$ , the directed path from  $r_{i'}$  to  $y$  on the red hand of the forward clock, and the directed path from  $y$  to  $r'_{i'}$  on the red hand of the reverse clock. Thus, the set  $X$  intersects every directed odd cycle of Type “Double Red”. Symmetrically, since  $a_i, a'_i \in X$ , we deduce that  $X$  also intersects every directed odd cycle of Type “Double Blue”. This concludes the proof of the reverse direction.  $\square$

**Lemma 3.7.** *Let  $\tilde{C}$  be a double clock. The weight of a set  $X \subseteq V(\tilde{C})$  that is a directed odd cycle transversal of  $\tilde{C}$  but does not cut  $\tilde{C}$  precisely is at least 70.*

*Proof.* Since  $X$  does not cut  $\tilde{C}$  precisely, Lemma 3.6 implies that either  $w(X) < 60$  or  $w(X) > 60$ . However, Lemmata 3.2 and 3.3 imply that the weight of the intersection of  $X$  with the forward clock is either 30 or at least 40, and Lemmata 3.4 and 3.5 imply that the weight of the intersection of  $X$  with the reverse clock is either 30 or at least 40. Furthermore, since  $y$  is the only vertex that the forward and reverse clocks have in common and its weight is  $2k + 1 > 70$ , we conclude that  $w(X) \geq 70$ .  $\square$

To analyze structures that combine several double clocks, we need to strengthen the reverse direction of Lemma 3.6. More precisely, we need to derive additional properties of a double clock from which we remove a set that cuts it precisely (in addition to the claim that this graph excludes directed odd cycles). For this purpose, we first introduce the following definition, which breaks a double clock into three “pieces” (see Fig. 9).

**Definition 3.4.** *Let  $\tilde{C}$  be a double clock, and let  $X$  be a set that cuts  $\tilde{C}$  precisely. Then,  $\tilde{C}[X, x]$  denotes the subgraph of  $\tilde{C} \setminus X$  induced by the set of vertices that both can reach  $x$  and are reachable from  $x$ , and  $\tilde{C}[X, z]$  denotes the subgraph of  $\tilde{C} \setminus X$  induced by the set of vertices that both can reach  $z$  and are reachable from  $z$ . Moreover,  $\tilde{C}[X, y]$  denotes the subgraph of  $\tilde{C} \setminus X$  induced by the set of vertices that belong to neither  $\tilde{C}[X, x]$  nor  $\tilde{C}[X, z]$ .*

Notice that for a double clock  $\tilde{C}$ , the only two directed paths from  $x$  to  $y$  are the red and blue hands of the forward clock of  $\tilde{C}$ , the only two directed paths from  $y$  to  $z$  are the red and blue hands of the reverse clock of  $\tilde{C}$ , and all of the directed paths from  $x$  to  $z$  contain the vertex  $y$ . Moreover, to every vertex  $v$  on a hand of the forward clock, there exists exactly one directed path from  $x$  that avoids  $y$ . On the other hand, from the vertex  $v$ , note that  $x$  is reachable by using an arc to the face of the forward clock (in case of a vertex of weight  $2k + 1$ ) or an outgoing arc followed by an arc to the face of the forward clock (in case of a vertex of weight 10), after which we append either of the two paths from the vertex we have reached on the face of the forward clock to  $x$ . A symmetric claim holds in the context of  $z$ . In light of these observations, we identify each piece of Definition 3.4 as follows (see Fig. 9).

**Observation 3.6.** *Let  $\tilde{C}$  be a double clock, and let  $X = \{p_i, a_{n-i+1}, p'_i, a'_{n-i+1}, \hat{t}_{i-1, n-i}, \hat{t}'_{i, n-i+1}\}$  be a set that cuts  $\tilde{C}$  precisely. Then, the following three conditions are satisfied.*

1.  $V(\tilde{C}[X, x]) = \hat{R} \cup \hat{B} \cup (\hat{T} \setminus \{\hat{t}_{i-1, n-i}\}) \cup \{x\} \cup \{p_{i'} \in P : i' < i\} \cup \{r_{i'} \in R : i' < i\} \cup \{a_{i'} \in A : i' < n - i + 1\} \cup \{b_{i'} \in B : i' < n - i + 1\}$ .
2.  $V(\tilde{C}[X, y]) = \{y\} \cup \{p_{i'} \in P : i' > i\} \cup \{r_{i'} \in R : i' \geq i\} \cup \{a_{i'} \in A : i' > n - i + 1\} \cup \{b_{i'} \in B : i' \geq n - i + 1\} \cup \{p'_{i'} \in P' : i' < i\} \cup \{r'_{i'} \in R' : i' < i\} \cup \{a'_{i'} \in A' : i' < n - i + 1\} \cup \{b'_{i'} \in B' : i' < n - i + 1\}$ .
3.  $V(\tilde{C}[X, z]) = \hat{R}' \cup \hat{B}' \cup (\hat{T}' \setminus \{\hat{t}'_{i, n-i+1}\}) \cup \{z\} \cup \{p'_{i'} \in P' : i' > i\} \cup \{r'_{i'} \in R' : i' \geq i\} \cup \{a'_{i'} \in A' : i' > n - i + 1\} \cup \{b'_{i'} \in B' : i' \geq n - i + 1\}$ .

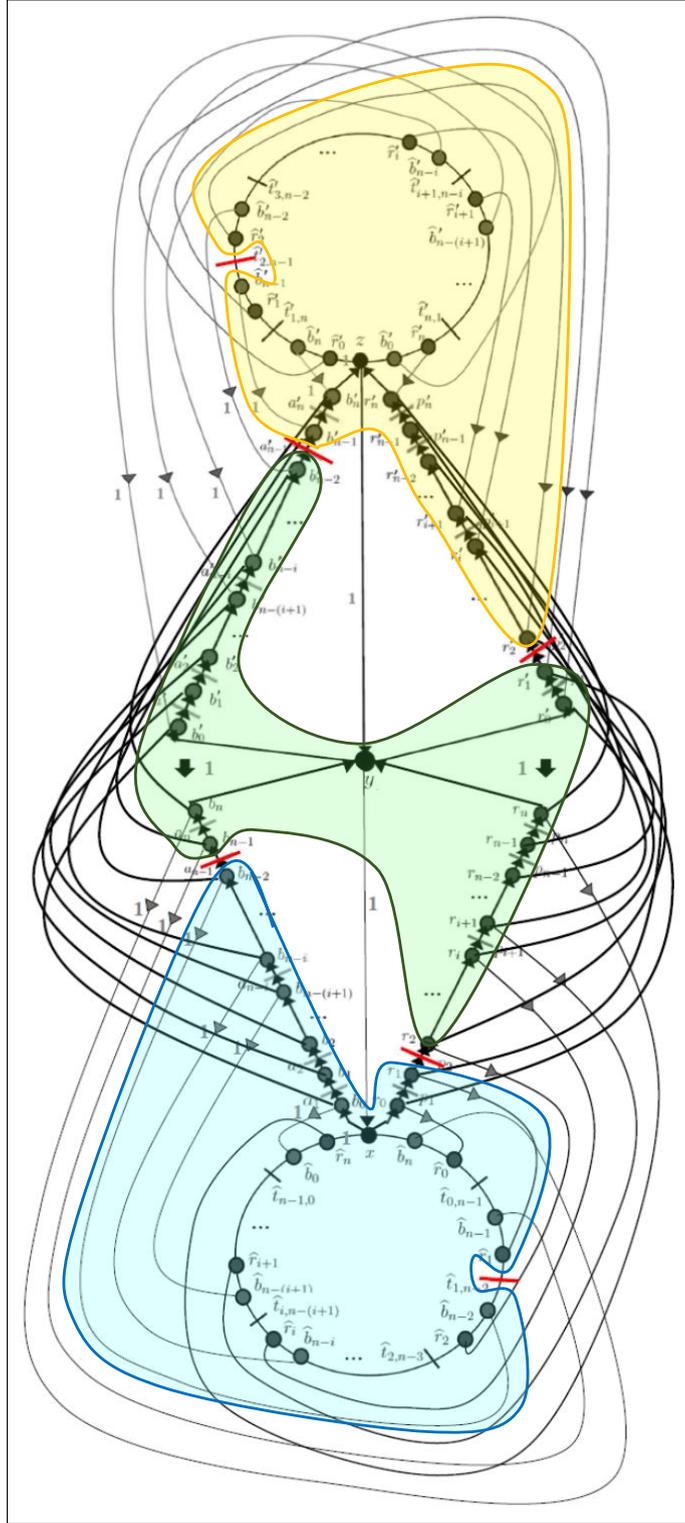


Figure 9: The components  $\tilde{C}[X, x]$  (blue),  $\tilde{C}[X, y]$  (green) and  $\tilde{C}[X, z]$  (yellow), where  $X$  is the set of red vertices (see Definition 3.4).

Furthermore, we notice that the three pieces are locally “isolated” in a double clock. Here, isolation means that there does not exist a vertex in one piece and a vertex in another piece such that the first vertex can reach the second one and vice versa. More precisely, since  $\tilde{C}[X, x]$

and  $\tilde{C}[X, z]$  are strongly connected directed graphs, while  $(z, y), (y, x) \in A(\tilde{C} \setminus X)$  and all of the directed paths of  $\tilde{C}$  from  $x$  to  $z$  contain  $y$ , we directly derive the following observation.

**Observation 3.7.** *Let  $\tilde{C}$  be a double clock, and let  $X$  be a set that cuts  $\tilde{C}$  precisely. Then, the following two conditions are satisfied.*

1. *There do not exist vertices  $u \in V(\tilde{C}[X, x])$  and  $v \in V(\tilde{C}[X, y]) \cup V(\tilde{C}[X, z])$  such that there exists a directed path from  $u$  to  $v$  in  $\tilde{C} \setminus X$ .*
2. *There do not exist vertices  $u \in V(\tilde{C}[X, y])$  and  $v \in V(\tilde{C}[X, z])$  such that there exists a directed path from  $u$  to  $v$  in  $\tilde{C} \setminus X$ .*

We also need to internally analyze each piece separately. To this end, we introduce one additional definition.

**Definition 3.5.** *Let  $D$  be a directed graph, and let  $\ell : A(D) \rightarrow \{0, 1\}$ . We say that a function  $f : V(D) \rightarrow \{\mathbf{b}, \mathbf{w}\}$  is  $\ell$ -consistent for  $D$  if for all  $(u, v) \in A(D)$ , it holds that  $\ell(u, v) = 0$  if and only if  $f(u) = f(v)$ .*

When the graph  $D$  is clear from context, we simply write  $\ell$ -consistent rather than  $\ell$ -consistent for  $D$ . First, we note the following simple observation, which hints at the relevance of Definition 3.5 to A-DOCT.

**Observation 3.8.** *Let  $D$  be a directed graph, and let  $\ell : A(D) \rightarrow \{0, 1\}$ . If there exists an  $\ell$ -consistent function for  $D$ , then  $D$  does not contain a directed odd cycle.*

Let us now derive another simple implication of Definition 3.5.

**Lemma 3.8.** *Let  $D$  be a directed graph,  $\ell : A(D) \rightarrow \{0, 1\}$ , and  $D'$  be some subgraph of  $D$  whose underlying undirected graph is connected and which contains only 0-labeled arcs. Then, if  $D$  admits an  $\ell$ -consistent function, then  $D$  also admits an  $\ell$ -consistent function  $f$  such that for all  $v \in V(D')$ , it holds that  $f(v) = \mathbf{b}$ .*

*Proof.* Suppose that  $D$  admits an  $\ell$ -consistent function  $\hat{f}$ . Then, define  $\hat{f}' : V(D) \rightarrow \{\mathbf{b}, \mathbf{w}\}$  as follows. For all  $v \in V(D)$ , it holds that  $\hat{f}'(v) = \mathbf{b}$  if and only if  $\hat{f}(v) = \mathbf{w}$ . Note that  $\hat{f}'$  is also  $\ell$ -consistent for  $D$ . By Definition 3.5, for all  $u, v \in V(D')$ , we have that  $\hat{f}(u) = \hat{f}(v)$  and  $\hat{f}'(u) = \hat{f}'(v)$ . Thus, if for all  $v \in V(D')$ , it holds that  $\hat{f}(v) = \mathbf{b}$ , then  $\hat{f}$  is a function  $f$  as stated in the lemma, and otherwise  $\hat{f}'$  is such a function  $f$ .  $\square$

We proceed by showing that for (arc-labeled) strongly connected directed graphs, we can easily find a consistent function.

**Lemma 3.9.** *Let  $D$  be a strongly connected directed graph, and let  $\ell : A(D) \rightarrow \{0, 1\}$ . If  $D$  does not contain a directed odd cycle, then  $D$  admits a function  $f$  that is  $\ell$ -consistent.*

*Proof.* Suppose that  $D$  does not contain a directed odd cycle. Let  $\hat{D}$  be the directed graph obtained from  $D$  by subdividing every 0-labeled arc once. That is, the graph  $\hat{D}$  is obtained from  $D$  by replacing every arc  $a = (u, v) \in A(D)$  that is labeled 0 by a new vertex  $w_a$  and the arcs  $(u, w_a)$  and  $(w_a, v)$ . Let  $\tilde{G}$  be the underlying undirected graph of  $\hat{D}$ . Note that  $V(D) \subseteq V(\hat{D}) = V(\tilde{G})$ . By Proposition 2.1,  $\tilde{G}$  is a bipartite graph. Then, there exists a bipartition  $(X, Y)$  of the vertex set of  $\tilde{G}$ . Define a function  $f : V(D) \rightarrow \{\mathbf{b}, \mathbf{w}\}$  as follows. For all  $v \in V(D)$ , it holds that  $f(v) = \mathbf{b}$  if and only if  $v \in X$ .

We claim that  $f$  is  $\ell$ -consistent for  $D$ . First, note that for every arc  $(u, v) \in A(D)$  that is labeled 1, the edge  $\{u, v\}$  belongs to  $E(\tilde{G})$ . Thus, since  $(X, Y)$  is a bipartition of  $\tilde{G}$ , it holds that either both  $u \in X$  and  $v \in Y$  or both  $v \in X$  and  $u \in Y$ . In either case, we have that  $f(u) \neq f(v)$ . Now, let  $a = (u, v)$  be some arc of  $D$  that is labeled 0. Then,  $\{u, w_a\}, \{w_a, v\} \in E(\tilde{G})$ . Thus, since  $(X, Y)$  is a bipartition of  $\tilde{G}$ , it holds that either  $u, v \in X$  or  $u, v \in Y$ . In either case, we have that  $f(u) = f(v)$ . This concludes the proof of the lemma.  $\square$

Finally, we are ready to present the last property of a double clock relevant to our work.

**Lemma 3.10.** *Let  $\tilde{C}$  be a double clock, and let  $X$  be a set that cuts  $\tilde{C}$  precisely. Then, the following three conditions are satisfied.*

1. *There exists an  $\ell$ -consistent function  $f_x$  for  $\tilde{C}[X, x]$  such that  $f_x(x) = \mathbf{b}$  and for every vertex  $v$  of  $\tilde{C}[X, x]$  that does not belong to the face of the forward clock, it holds that  $f_x(v) = \mathbf{b}$ .*
2. *The function that assigns  $\mathbf{b}$  to every vertex of  $\tilde{C}[X, y]$  is  $\ell$ -consistent for  $\tilde{C}[X, y]$ .*
3. *There exists an  $\ell$ -consistent function  $f_z$  for  $\tilde{C}[X, z]$  such that  $f_z(z) = \mathbf{b}$  and for every vertex  $v$  of  $\tilde{C}[X, z]$  that does not belong to the face of the reverse clock, it holds that  $f_z(v) = \mathbf{b}$ .*

*Proof.* By the definitions of  $\tilde{C}[X, x]$  and  $\tilde{C}[X, z]$ , we have that  $\tilde{C}[X, x]$  and  $\tilde{C}[X, z]$  are strongly connected directed graphs. Moreover, by Lemma 3.6,  $\tilde{C} \setminus X$  does not contain a directed odd cycle, and therefore both  $\tilde{C}[X, x]$  and  $\tilde{C}[X, z]$  do not contain a directed odd cycle. Thus, by Lemma 3.9, both  $\tilde{C}[X, x]$  and  $\tilde{C}[X, z]$  admit  $\ell$ -consistent functions. By Observation 3.6 and the definition of a double clock, we have that  $\tilde{C}[X, x]$  contains a subgraph whose vertex set consists of  $x$  as well as every vertex of  $\tilde{C}[X, x]$  that does not belong to the face of the forward clock, whose underlying undirected graph is connected, and which consists only of 0-labeled arcs. Symmetrically, we have that  $\tilde{C}[X, z]$  contains a subgraph whose vertex set consists of  $z$  as well as every vertex of  $\tilde{C}[X, z]$  that does not belong to the face of the reverse clock, whose underlying undirected graph is connected, and which consists only of 0-labeled arcs. Hence, by Lemma 3.8, we have that Conditions 1 and 3 are satisfied.

Finally, by Observation 3.6 and the definition of a double clock, we note that all of the arcs of  $\tilde{C}[X, y]$  are labeled 0. Thus, it is clear that Condition 2 is satisfied as well.  $\square$

### 3.4 The Synchronization Gadget

Let  $n, k \in \mathbb{N}$  such that  $k \geq 100$ , and let  $I \subseteq [n] \times [n]$  be a set of pairs of indices. Here, we define an  $(n, k, I)$ -synchronizer. Since  $n$  and  $k$  would be clear from context, we simply write  $I$ -synchronizer rather than  $(n, k, I)$ -synchronizer. When  $I$  is also clear from context (or immaterial), we omit it as well.

**Structure.** The *hands* of a synchronizer  $S$  are four red directed paths,  $H, H', \tilde{H}$  and  $\tilde{H}'$ . The vertex set of  $H$  is the union of two pairwise disjoint sets,  $R = \{r_i : i \in [n]_0\}$  (red) and  $P = \{p_i : i \in [n]\}$  (pink). For all  $i \in [n]$ , we denote  $\text{pre}(p_i) = r_{i-1}$  and  $\text{post}(p_i) = r_i$ . The arc set of  $H$  is  $\{(\text{pre}(p_i), p_i) : i \in [n]\} \cup \{(p_i, \text{post}(p_i)) : i \in [n]\}$  (see Fig. 10). The path  $\tilde{H}$  is defined as the path  $H$  where we use tilde notation to specify vertices. Similarly,  $H'$  and  $\tilde{H}'$  are defined as the path  $H$  and  $\tilde{H}$ , respectively, where we further use prime notation to specify vertices. The weight of each vertex on these paths is 10, and the label of each arc on these paths is 0. Now, to obtain the *frame* of  $S$ , we add three vertices,  $x, y$  and  $z$ , each of weight  $2k + 1$ . Moreover, we add the arcs  $(x, r_0), (x, \tilde{r}_0), (r_n, y), (\tilde{r}_n, y), (y, r'_0), (y, \tilde{r}'_0), (r'_n, z)$  and  $(\tilde{r}'_n, z)$ . The label of each of these arcs is 0. For the sake of clarity of illustrations, the vertex  $y$  is drawn twice (see Fig. 10).

Next, we define the *interior* of  $S$  (see Fig. 10). Roughly speaking, this part is a grid where each vertex has either a very high weight or a very low weight, depending on whether or not the pair of indices that the vertex represents belongs to  $I$ . Formally, the interior of  $S$  is the graph  $G$  on the vertex set  $\{g_{i,j} : i, j \in [n]\}$  and the arc set  $\{(g_{i+1,j}, g_{i,j}) : i \in [n-1], j \in [n]\} \cup \{(g_{i,j+1}, g_{i,j}) : i \in [n], j \in [n-1]\}$ . The label of each of the arcs is 0. Moreover, for all  $i, j \in [n]$ , the weight of  $g_{i,j}$  is 1 if  $(i, j) \in I$  and  $2k + 1$  otherwise.

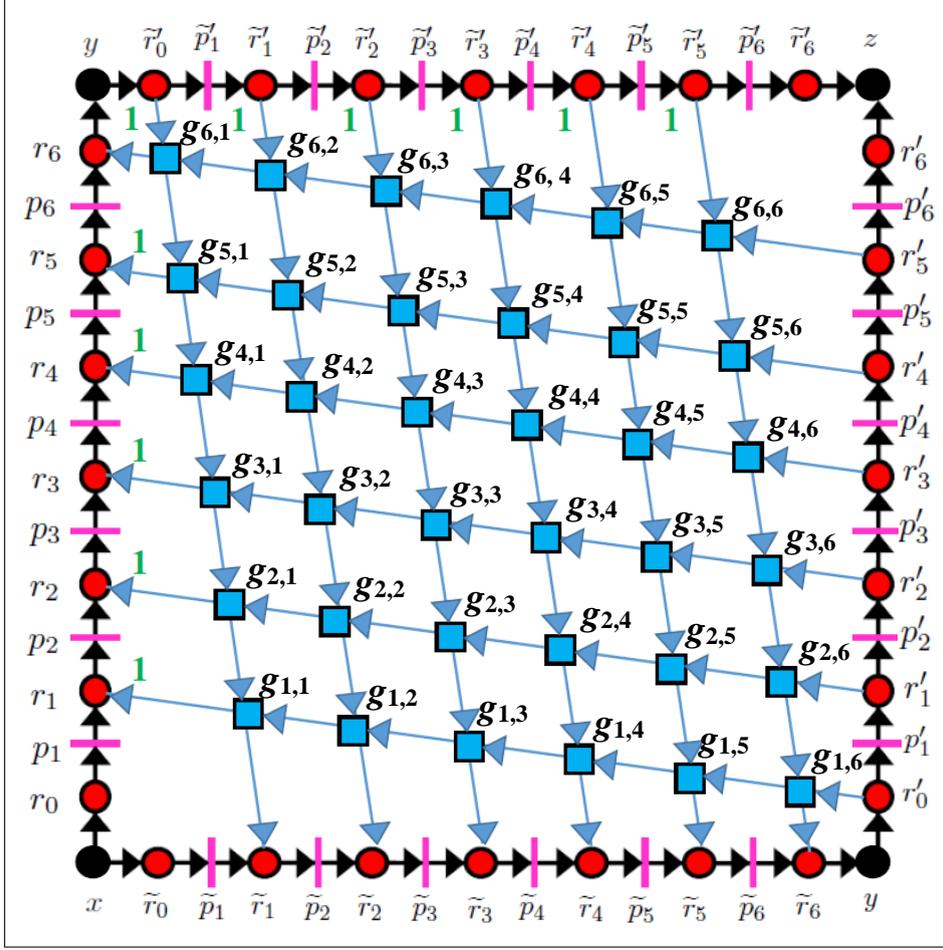


Figure 10: A synchronizer where  $n = 6$ . The arcs labeled 1 are marked by a green ‘1’. The weight of vertices marked by circles is  $2k + 1$ , the weight of vertices marked by lines is 10, and the weight of each vertex marked by a square is either  $2k + 1$  or 1.

Finally, we attach the frame of  $S$  to the interior of  $S$  (see Fig. 10). To this end, for all  $i \in [n]$ , we add two arcs labeled 1:  $(g_{i,1}, \text{post}(p_i))$  and  $(\text{pre}(\tilde{p}'_i), g_{n,i})$ . Moreover, for all  $i \in [n]$ , we add two arcs labeled 0:  $(g_{1,i}, \text{post}(\tilde{p}_i))$  and  $(\text{pre}(p'_i), g_{i,n})$ . When the synchronizer  $S$  is not clear from context, we add the notation  $(S)$  to an element (vertex set or vertex) of the synchronizer.

**Properties.** By the definition of a synchronizer, we first directly identify which directed odd cycles are present in such a gadget.

**Observation 3.9.** *Let  $S$  be a synchronizer. The set of directed odd cycles of  $S$  is the union of the following sets.*

- **Horizontal Match:** For all  $i \in [n]$ , this set contains the direct odd cycle consisting of the directed path from  $y$  to  $\text{pre}(p'_i)$  on  $H'$ , the (unique) directed path from  $\text{pre}(p'_i)$  to  $\text{post}(p_i)$  on the interior, and the directed path from  $\text{post}(p_i)$  to  $y$  on  $H$ .
- **Horizontal Mismatch:** For all  $i, j \in [n]$  such that  $j < i$ , this set contains every direct odd cycle consisting of the directed path from  $y$  to  $\text{pre}(p'_i)$  on  $H'$ , some directed path from  $\text{pre}(p'_i)$  to  $\text{post}(p_j)$  on the interior, and the directed path from  $\text{post}(p_j)$  to  $y$  on  $H$ .
- **Vertical Match:** For all  $i \in [n]$ , this set contains the direct odd cycle consisting of the directed path from  $y$  to  $\text{pre}(\tilde{p}'_i)$  on  $\tilde{H}'$ , the (unique) directed path from  $\text{pre}(\tilde{p}'_i)$  to  $\text{post}(\tilde{p}_i)$  on the interior, and the directed path from  $\text{post}(\tilde{p}_i)$  to  $y$  on  $\tilde{H}$ .

- **Vertical Mismatch:** For all  $i, j \in [n]$  such that  $j < i$ , this set contains every directed odd cycle consisting of the directed path from  $y$  to  $\text{pre}(\tilde{p}'_i)$  on  $\tilde{H}'$ , some directed path from  $\text{pre}(\tilde{p}'_i)$  to  $\text{post}(\tilde{p}_j)$  on the interior, and the directed path from  $\text{post}(\tilde{p}_j)$  to  $y$  on  $\tilde{H}$ .

We proceed to derive properties of “cuts” of a synchronizer. To this end, we again first need to define the kind of sets using which we would like to “cut” synchronizers.

**Definition 3.6.** Let  $S$  be an  $I$ -synchronizer. We say that a set  $X \subseteq V(S)$  cuts  $S$  precisely if there exist  $i, j \in [n]$  such that  $X = \{p_i, p'_i, \tilde{p}_j, \tilde{p}'_j, g_{i,j}\}$  and  $(i, j) \in I$ .

**Definition 3.7.** Let  $S$  be an  $I$ -synchronizer. We say that a set  $X \subseteq V(S)$  cuts  $S$  roughly if  $X$  does not cut  $S$  precisely and there exist  $i, j \in [n]$  such that  $\{p_i, p'_i, \tilde{p}_j, \tilde{p}'_j\} \subseteq X$ .

We are now ready to present desired properties of “cuts” of a synchronizer. Unlike the cases of clocks, here we only analyze cuts of the forms presented in Definitions 3.6 and 3.7. The first property follows directly from Definition 3.6.

**Observation 3.10.** The weight of a set that cuts a synchronizer precisely is exactly 41. In particular, the weight of the intersection of this set with the interior is exactly 1.

Let us now argue that all directed odd cycles are intersected.

**Lemma 3.11.** Let  $S$  be a synchronizer, and let  $X$  be a set that cuts  $S$  precisely. Then,  $S \setminus X$  does not contain a directed odd cycle.

*Proof.* Since  $X$  cuts  $S$  precisely, there exist  $i, j \in [n]$  such that  $X = \{p_i, p'_i, \tilde{p}_j, \tilde{p}'_j, g_{i,j}\}$ . Since  $g_{i,j} \in X$ , it holds that  $X$  intersects the directed odd cycle of Type “Horizontal Match” that consists of the directed path from  $y$  to  $\text{pre}(p'_i)$  on  $H'$ , the directed path from  $\text{pre}(p'_i)$  to  $\text{post}(p_i)$  on the interior, and the directed path from  $\text{post}(p_i)$  to  $y$  on  $H$  (see Observation 3.9). Moreover, since  $p_i, p'_i \in X$ , it holds that  $X$  intersects all of the remaining directed odd cycles of Types “Horizontal Match” and “Horizontal Mismatch”. Symmetrically, since  $g_{i,j} \in X$ , it holds that  $X$  intersects the directed odd cycle of Type “Vertical Match” that consists of the directed path from  $y$  to  $\text{pre}(\tilde{p}'_j)$  on  $\tilde{H}'$ , the directed path from  $\text{pre}(\tilde{p}'_j)$  to  $\text{post}(\tilde{p}_j)$  on the interior, and the directed path from  $\text{post}(\tilde{p}_j)$  to  $y$  on  $\tilde{H}$ . Finally, since  $\tilde{p}_j, \tilde{p}'_j \in X$ , it holds that  $X$  intersects all of the remaining directed odd cycles of Types “Vertical Match” and “Vertical Mismatch”. This concludes the proof of the lemma.  $\square$

Next, we analyze the weight of rough cuts.

**Lemma 3.12.** Let  $S$  be a synchronizer, and let  $X$  be a set that cuts  $S$  roughly. Then,  $w(X) \geq 42$ .

*Proof.* If  $X$  contains at least five vertices of the frame, then  $w(X) \geq 50$  since the weight of each vertex of the frame is at least 10. Thus, we next assume that  $X$  contains exactly four vertices of the frame. Since  $X$  cuts  $S$  roughly, the set of these vertices can be denoted by  $Y = \{p_i, p'_i, \tilde{p}_j, \tilde{p}'_j\}$  for some  $i, j \in [n]$ . Let  $C$  be the directed odd cycle of Type “Horizontal Match” consisting of the directed path from  $y$  to  $\text{pre}(p'_i)$  on  $H'$ , the directed path from  $\text{pre}(p'_i)$  to  $\text{post}(p_i)$  on the interior, and the directed path from  $\text{post}(p_i)$  to  $y$  on  $H$  (see Observation 3.9). Moreover, let  $C'$  be the directed odd cycle of Type “Vertical Match” consisting of the directed path from  $y$  to  $\text{pre}(\tilde{p}'_j)$  on  $\tilde{H}'$ , the (unique) directed path from  $\text{pre}(\tilde{p}'_j)$  to  $\text{post}(\tilde{p}_j)$  on the interior, and the directed path from  $\text{post}(\tilde{p}_j)$  to  $y$  on  $\tilde{H}$ . Note that besides vertices of the frame, the only vertex that both  $C$  and  $C'$  have in common is  $g_{i,j}$ . Since  $X$  must intersect both of these cycles but it cannot contain only  $g_{i,j}$  in addition to  $Y$  (since it does not cut  $S$  precisely), we deduce that  $X$  must contain at least two vertices of the interior. We thus conclude that  $w(X) \geq 42$ .  $\square$

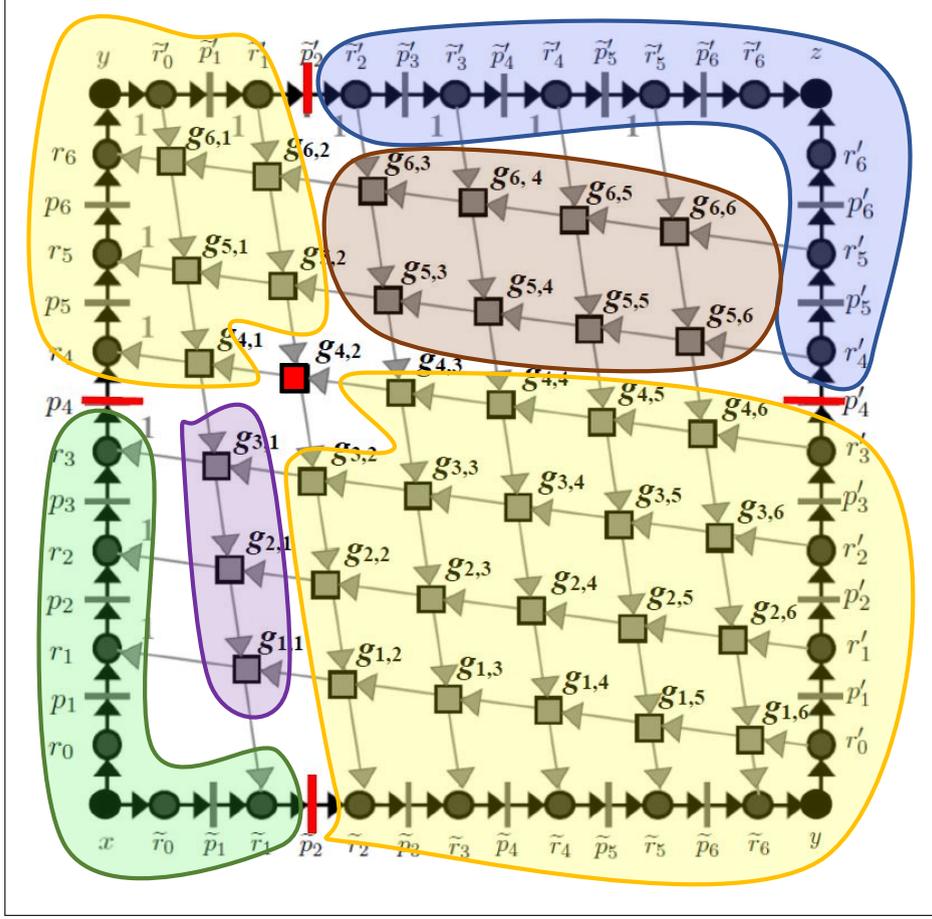


Figure 11: The components  $S[X, x]$  (green),  $S[X, y]$  (yellow),  $S[X, z]$  (blue),  $S[X, l_x]$  (purple) and  $S[X, l_z]$  (brown), where  $X$  is the set of red vertices (see Definition 3.8).

As in the case of the double clock, we need to strengthen Lemma 3.11. For this purpose, we introduce the following definition, which breaks a synchronizer into five “pieces” (see Fig. 11).

**Definition 3.8.** *Let  $S$  be a synchronizer, and let  $X$  be a set that cuts  $S$  precisely. Then,  $S[X, y]$  denotes the subgraph of  $S \setminus X$  induced by the set of vertices that both can reach  $y$  and are reachable from  $y$ . Moreover,  $S[X, x]$  denotes the subgraph of  $S \setminus X$  induced by the set of vertices reachable from  $x$ , and  $S[X, z]$  denotes the subgraph of  $S$  induced by the set of vertices that can reach  $z$ . Finally,  $S[X, l_x]$  denotes the subgraph of  $S \setminus X$  induced by the set of vertices outside  $S[X, x]$  that can reach a vertex of  $S[X, x]$  without using any vertex of  $S[X, y]$ , and  $S[X, l_z]$  denotes the subgraph of  $S \setminus X$  induced by the set of vertices outside  $S[X, z]$  that are reachable from a vertex of  $S[X, z]$  without using any vertex of  $S[X, y]$ .*

Notice that for a synchronizer  $S$ , the only two directed paths from  $x$  to  $y$  are those internally consisting of  $H$  and  $\tilde{H}$ , the only two directed paths from  $y$  to  $z$  are those internally consisting of  $H'$  and  $\tilde{H}'$ , and all of the directed paths from  $x$  to  $z$  contain the vertex  $y$ . Moreover, the vertex  $y$  and every vertex  $g_{i,j}$  of the interior are both contained in the two following (even) directed cycles, among other directed cycles, whose only common vertices are  $y$  and  $g_{i,j}$ : (i) the directed cycle consisting of the path from  $y$  to  $\text{pre}(p'_i)$  on  $H'$ , the (unique) directed path from  $\text{pre}(p'_i)$  to  $g_{i,j}$  on the interior, the (unique) directed path from  $g_{i,j}$  to  $\text{post}(\tilde{p}_j)$  on the interior, and the directed path from  $\text{post}(\tilde{p}_j)$  to  $y$  on  $\tilde{H}$ ; (ii) the directed cycle consisting of the path from  $y$  to  $\text{pre}(\tilde{p}'_j)$  on  $\tilde{H}'$ , the (unique) directed path from  $\text{pre}(\tilde{p}'_j)$  to  $g_{i,j}$  on the interior, the (unique)

directed path from  $g_{i,j}$  to  $\text{post}(p_i)$  on the interior, and the directed path from  $\text{post}(p_i)$  to  $y$  on  $H$ . In light of these observations, we identify each piece of Definition 3.8 as follows (see Fig. 11).

**Observation 3.11.** *Let  $S$  be a synchronizer, and let  $X = \{p_i, p'_i, \tilde{p}_j, \tilde{p}'_j, g_{i,j}\}$  be a set that cuts  $S$  precisely. Then, the following five conditions are satisfied.*

1.  $V(S[X, x]) = \{x\} \cup \{p_{i'} \in V(H) : i' < i\} \cup \{r_{i'} \in V(H) : i' < i\} \cup \{\tilde{p}_{j'} \in V(\tilde{H}) : j' < j\} \cup \{\tilde{r}_{j'} \in V(\tilde{H}) : j' < j\}$ .
2.  $V(S[X, y]) = \{y\} \cup \{p_{i'} \in V(H) : i' > i\} \cup \{r_{i'} \in V(H) : i' \geq i\} \cup \{\tilde{p}_{j'} \in V(\tilde{H}) : j' > j\} \cup \{\tilde{r}_{j'} \in V(\tilde{H}) : j' \geq j\} \cup \{p'_{i'} \in V(H') : i' < i\} \cup \{r'_{i'} \in V(H') : i' < i\} \cup \{\tilde{p}'_{j'} \in V(\tilde{H}') : j' < j\} \cup \{\tilde{r}'_{j'} \in V(\tilde{H}') : j' < j\} \cup (\{g_{i',j'} \in V(G) : i' \leq i, j' \geq j\} \cup \{g_{i',j'} \in V(G) : i' \geq i, j' \leq j\}) \setminus \{g_{i,j}\}$ .
3.  $V(S[X, z]) = \{z\} \cup \{p'_{i'} \in V(H') : i' > i\} \cup \{r'_{i'} \in V(H') : i' \geq i\} \cup \{\tilde{p}'_{j'} \in V(\tilde{H}') : j' > j\} \cup \{\tilde{r}'_{j'} \in V(\tilde{H}') : j' \geq j\}$ .
4.  $V(S[X, l_x]) = \{g_{i',j'} \in V(G) : i' < i, j' < j\}$ .
5.  $V(S[X, l_z]) = \{g_{i',j'} \in V(G) : i' > i, j' > j\}$ .

Furthermore, in light of Observation 3.11, we notice that the five pieces are locally “isolated” in a synchronizer as follows.

**Observation 3.12.** *Let  $S$  be a synchronizer, and let  $X$  be a set that cuts  $S$  precisely. Then, the following five conditions are satisfied.*

1. *There do not exist vertices  $u \in V(\tilde{C}[X, x])$  and  $v \in V(S[X, l_x]) \cup V(S[X, y]) \cup V(S[X, l_z]) \cup V(S[X, z])$  such that there exists a directed path from  $u$  to  $v$  in  $S \setminus X$ .*
2. *There do not exist vertices  $u \in V(\tilde{C}[X, l_x])$  and  $v \in V(S[X, y]) \cup V(S[X, l_z]) \cup V(S[X, z])$  such that there exists a directed path from  $u$  to  $v$  in  $S \setminus X$ .*
3. *There do not exist vertices  $u \in V(\tilde{C}[X, y])$  and  $v \in V(S[X, l_z]) \cup V(S[X, z])$  such that there exists a directed path from  $u$  to  $v$  in  $S \setminus X$ .*
4. *There do not exist vertices  $u \in V(\tilde{C}[X, l_z])$  and  $v \in V(S[X, z])$  such that there exists a directed path from  $u$  to  $v$  in  $S \setminus X$ .*

Finally, we need to internally analyze each piece separately.

**Lemma 3.13.** *Let  $S$  be a synchronizer, and let  $X$  be a set that cuts  $S$  precisely. Then, the following two conditions are satisfied.*

1. *For each of the graphs  $S[X, x], S[X, l_x], S[X, l_z]$  and  $S[X, z]$ , the function that assigns  $\mathbf{b}$  to every vertex of the graph is  $\ell$ -consistent.*
2. *There exists an  $\ell$ -consistent function  $f_y$  for  $S[X, y]$  such that for every vertex  $v$  of the frame that belongs to  $S[X, y]$ , it holds that  $f_y(v) = \mathbf{b}$ .*

*Proof.* By Observation 3.11 and the definition of a synchronizer, we note that all of the arcs of each of the graphs  $S[X, x], S[X, l_x], S[X, l_z]$  and  $S[X, z]$  are labeled 0. Thus, it is clear that Condition 1 is satisfied

Next, by its definition, note that  $S[X, y]$  is a strongly connected directed graph. Moreover, by Lemma 3.11,  $S \setminus X$  does not contain a directed odd cycle, and therefore  $S[X, y]$  does not contain a directed odd cycle. Thus, by Lemma 3.9,  $S[X, y]$  admits an  $\ell$ -consistent function. By

Observation 3.11 and the definition of a synchronizer, we have that  $S[X, y]$  contains a subgraph whose vertex set consists of all of the vertices of the frame that belong to  $S[X, y]$ , whose underlying undirected graph is connected, and which consists only of 0-labeled arcs. Hence, by Lemma 3.8, we have that Condition 2 is satisfied as well.  $\square$

### 3.5 Reduction

We are now ready to present the complete reduction from PSI to A-DOCT. For this purpose, let  $(H, G, col)$  be an instance PSI. We assume that  $|V(G)| \geq 100$ , else a solution can be found by brute force in polynomial time. If  $G$  contains an isolated vertex to which no vertex in  $H$  is mapped by  $col$ , then the input instance is a no-instance; otherwise, by removing all of the isolated vertices of  $G$  and the vertices of  $H$  that are mapped to them, we obtain an instance of PSI that is equivalent to  $(H, G, col)$ . Thus, we next assume that  $G$  does not contain isolated vertices. For all  $g \in V(G)$ , denote  $V^g = \{v \in V(H) : col(v) = g\}$ . We next assume that for all  $g, g' \in V(G)$ , it holds that  $|V^g| = |V^{g'}| = n$  (for the appropriate  $n$ ), else we can add isolated vertices to  $H$  to ensure that this condition holds. Then, for all  $g \in V^g$ , denote  $V^g = \{v_1^g, v_2^g, \dots, v_n^g\}$ . Let  $<$  be some arbitrary order on  $V(G)$ .

We construct an instance  $\mathbf{red}(H, G, col) = (D, k, \ell, w)$  of A-DOCT as follows. First, we set  $k = 60|V(G)| + |E(G)|$ . Next, we turn to construct  $(D, k, \ell, w)$ . For every  $g \in V(G)$ , we insert one  $(n, k)$ -double clock  $\tilde{C}^g$ . For every edge  $e = \{g, g'\} \in E(G)$  where  $g < g'$ , we insert one  $(n, k, I^e)$ -synchronizer  $S^e$  where  $I^e = \{(i, j) : \{v_i^g, v_j^{g'}\} \in E(H)\}$ . We identify the vertices  $x, y$  and  $z$  of all double clocks and synchronizers. That is, we now have a single vertex called  $x$ , a single vertex called  $y$  and a single vertex called  $z$ . Finally, for every edge  $e = \{g, g'\} \in E(G)$  where  $g < g'$ , we identify the red hand of the forward clock of  $\tilde{C}^g$  with the hand  $H$  of  $S^e$ , the red hand of the reverse clock of  $\tilde{C}^g$  with the hand  $H'$  of  $S^e$ , the red hand of the forward clock of  $\tilde{C}^{g'}$  with the hand  $\tilde{H}$  of  $S^e$ , and the red hand of the reverse clock of  $\tilde{C}^{g'}$  with the hand  $\tilde{H}'$  of  $S^e$ . Here, by identifying two directed paths of the same number  $2n + 1$  of vertices, we mean that for all  $i \in [2n + 1]$ , we identify the  $i$ th vertex on one path with the  $i$ th vertex on the other path. Consequently, for all  $i \in [2n]$ , we also identify the  $i$ th arc on one path with the  $i$ th arc on the other path. We remark that next, when we refer to an element of a specific double clock or a specific synchronizer, we would refer to the new unified vertex. For example, given an edge  $e = \{g, g'\} \in E(G)$  where  $g < g'$ , we have that  $r_3(C^g) = r_3(S^e)$  and  $r'_5(C^{g'}) = \tilde{r}'_5(S^e)$ . This completes the description of the reduction.

### 3.6 Correctness

It remains to derive the correctness of Theorem 1. To this end, first note that since  $|V(G)| \geq 100$  and  $G$  does not contain isolated vertices, we have the following observation.

**Observation 3.13.** *Let  $(H, G, col)$  be an instance of PSI. Then, for  $(D, k, \ell, w) = \mathbf{red}(H, G, col)$ , it holds that  $100 \leq k \leq 121|E(G)|$ .*

Clearly, we also have the following observation.

**Observation 3.14.** *Let  $(H, G, col)$  be an instance of PSI. Then, the instance  $\mathbf{red}(H, G, col)$  can be constructed in polynomial time.*

To verify the correctness of the reduction, we first prove the forward direction, summarized in the following lemmata.

**Lemma 3.14.** *Let  $(H, G, col)$  be a yes-instance of PSI. Then,  $(D, k, \ell, w) = \mathbf{red}(H, G, col)$  is a yes-instance of A-DOCT.*

*Proof.* Since  $(H, G, \text{col})$  be a yes-instance of PSI, there exists a colorful mapping of  $G$  into  $H$ . That is, there exists an injective function  $\varphi : V(G) \rightarrow V(H)$  such that for every  $g \in V(G)$ ,  $\text{col}(\varphi(g)) = g$ , and for every  $\{g, g'\} \in E(G)$ ,  $\{\varphi(g), \varphi(g')\} \in E(H)$ . For all  $g \in V(G)$ , let  $i(g)$  denote the index  $i \in [n]$  such that  $\varphi(g) = v_i^g$ . Now, we define the following sets.

- For all  $g \in V(G)$ : We define  $X^g = \{p_i(\tilde{C}^g), a_{n-i+1}(\tilde{C}^g), p'_i(\tilde{C}^g), a'_{n-i+1}(\tilde{C}^g), \hat{t}_{i-1, n-i}(\tilde{C}^g), \hat{t}'_{i, n-i+1}(\tilde{C}^g)\}$  where  $i = i(g)$ . Note that  $X^g$  cuts the double clock  $\tilde{C}^g$  precisely.
- For all  $e = \{g, g'\} \in E(G)$  where  $g < g'$ : We define  $X^e = \{p_i(S^e), p'_i(S^e), \tilde{p}_j(S^e), \tilde{p}'_j(S^e), g_{i,j}(S^e)\}$  where  $i = i(g)$  and  $j = i(g')$ . Since  $\{\varphi(g), \varphi(g')\} \in E(H)$ , it holds that  $(i, j) \in I$ . Thus, we have that  $X^e$  cuts the synchronizer  $S^e$  precisely.

Accordingly, define  $X = (\bigcup_{g \in V(G)} X^g) \cup (\bigcup_{e \in E(G)} X^e)$ . By the definition of  $\text{red}(H, G, \text{col})$  and Observation 3.10, it holds that  $w(X) = 60|V(G)| + |E(G)|$ . We claim that  $X$  is a directed odd cycle transversal of  $D$ , which would imply that  $(D, k, \ell, w)$  is a yes-instance of A-DOCT. To this end, we also define the following sets.

- $R_x = \bigcup_{g \in V(G)} V(\tilde{C}^g[X^g, x])$ . Note that  $\bigcup_{e \in E(G)} V(S^e[X^e, x]) \subseteq R_x$ .
- $R_z = \bigcup_{g \in V(G)} V(\tilde{C}^g[X^g, z])$ . Note that  $\bigcup_{e \in E(G)} V(S^e[X^e, z]) \subseteq R_z$ .
- $R_y = \bigcup_{e \in E(G)} V(S^e[X^e, y])$ . Note that  $\bigcup_{g \in V(G)} V(\tilde{C}^g[X^g, y]) \subseteq R_y$ .
- $R_x^l = \bigcup_{e \in E(G)} V(S^e[X^e, l_x])$ .
- $R_z^l = \bigcup_{e \in E(G)} V(S^e[X^e, l_z])$ .

Denote  $\mathcal{R} = \{R_x, R_z, R_y, R_x^l, R_z^l\}$ . Note that the sets in  $\mathcal{R}$  are pairwise disjoint. By Observations 3.7 and 3.12, there does not exist a directed odd cycle whose vertex set intersects more than one set from  $\mathcal{R}$ . Moreover, by Lemmata 3.10 and 3.13, for every  $R \in \mathcal{R}$ , there exists a function that is  $\ell$ -consistent for  $D[R \setminus X]$ . By Observation 3.8, we deduce that for every  $R \in \mathcal{R}$ ,  $D[R \setminus X]$  does not contain a directed odd cycle. We thus derive that  $X$  is a directed odd cycle transversal of  $D$ . This concludes the proof of the lemma.  $\square$

We next prove a strengthened version of the reverse direction, which would be necessary to derive our inapproximability result.

**Lemma 3.15.** *Fix  $\epsilon \geq 0$ . There exists  $\delta = \delta(\epsilon) \geq 0$ , where if  $\epsilon > 0$  then  $\delta > 0$ , such that the following condition holds.*

- *Given an instance  $(H, G, \text{col})$  of PSI, if  $\text{red}(H, G, \text{col})$  admits a  $(1 + \delta)$ -approximate solution, then there exists a colorful mapping of a subgraph  $G'$  of  $G$  into  $H$  such that  $G'$  contains at least  $(1 - \epsilon)|E(G)|$  edges.*

*Proof.* Let  $\delta \leq 1$  be determined later. Suppose that  $(H, G, \text{col})$  is an instance of PSI such that  $\text{red}(H, G, \text{col})$  admits a  $(1 + \delta)$ -approximate solution  $X$ . We need to show that there exists a subgraph  $G'$  of  $G$  with at least  $(1 - \epsilon)|E(G)|$  edges such that there exists a colorful mapping of  $G'$  into  $H$ . Since  $X$  is a  $(1 + \delta)$ -approximate solution, it holds that  $w(X) \leq (1 + \delta)k = (1 + \delta)(60|V(G)| + |E(G)|)$ . Since  $\delta \leq 1$ , it holds that  $X$  excludes every vertex of  $D$  weight  $2k+1$ . In particular,  $X$  does not contain any vertex that is present in more than one double clock. Let  $\mathcal{C}^*$  be the set of every double clock whose intersection with  $X$  cuts it precisely. Moreover, let  $\mathcal{S}$  be the set of every synchronizer such that each of its four hands has been identified with a hand of a double clock from  $\mathcal{C}^*$ . Moreover, let  $\mathcal{S}^*$  be the set of every synchronizer in  $\mathcal{S}$  whose

intersection with  $X$  cuts it precisely. Denote  $c^* = |\mathcal{C}^*|$ ,  $s = |\mathcal{S}|$  and  $s^* = |\mathcal{S}^*|$ . Then, by Lemmata 3.6, 3.7, 3.10 and 3.12, it holds that

$$w(X) \geq 60c^* + 70(|V(G)| - c^*) + 2(s - s^*) + s^* = 70|V(G)| - 10c^* + 2s - s^*.$$

Hence, we have that  $70|V(G)| - 10c^* + 2s - s^* \leq (1 + \delta)(60|V(G)| + |E(G)|)$ . Thus,

$$2s + 70|V(G)| - (1 + \delta)(60|V(G)| + |E(G)|) \leq 10c^* + s^*.$$

Since  $G$  is a graph of maximum degree 3, it holds that  $|E(G)| - s \leq 3(|V(G)| - c^*)$ . That is,  $s \geq |E(G)| - 3(|V(G)| - c^*)$ . Then, by the inequality above,

$$64|V(G)| + 2|E(G)| - (1 + \delta)(60|V(G)| + |E(G)|) \leq 4c^* + s^*.$$

That is,

$$(4 - 60\delta)|V(G)| + (1 - \delta)|E(G)| \leq 4c^* + s^*.$$

Since  $c^* \leq |V(G)|$ , it holds that  $s^* \geq (1 - \delta)|E(G)| - 60\delta|V(G)|$ . Then, since  $G$  does not contain isolated vertices,  $|V(G)| \leq 2|E(G)|$ . Thus,  $s^* \geq (1 - 121\delta)|E(G)|$ . Fix  $\delta = \min\{\epsilon/121, 1\}$ . Note that if  $\epsilon > 0$  then  $\delta > 0$ . Overall, it holds that

$$s^* \geq (1 - \epsilon)|E(G)|.$$

Define  $E^* = \{e \in E(G) : S^e \in \mathcal{S}^*\}$ , and let  $V^*$  be the set of every vertex in  $V(G)$  that is incident in  $G$  to at least one edge in  $E^*$ . Note that  $|E^*| = s^*$ . Hence, to conclude that the lemma is correct, it is sufficient to show that  $(H, G^*, col)$  is a yes-instance. That is, we need to show that there exists an injection  $\varphi : V(G^*) \rightarrow V(H)$  such that for every  $g \in V(G^*)$ ,  $col(\varphi(g)) = g$ , and for every  $\{g, g'\} \in E^*$ ,  $\{\varphi(g), \varphi(g')\} \in E(H)$ . To this end, we define an injection  $\psi : V(G^*) \rightarrow V(H)$  as follows. For all  $g \in V(G^*)$ , since  $\tilde{C}^g \in \mathcal{C}^*$ , we can define  $i(g) \in [n]$  as the index such that the intersection of the approximate solution  $X$  with  $\tilde{C}^g$  is equal to  $\{p_{i(g)}, a_{n-i(g)+1}, p'_{i(g)}, a'_{n-i(g)+1}, \hat{t}_{i(g)-1, n-i(g)}, \hat{t}'_{i(g), n-i(g)+1}\}$ . Then, for all  $g \in V(G^*)$ , set  $\psi(g) = v_{i(g)}^g$ , and note that since  $v_{i(g)}^g \in V^g$ , it holds that  $col(\psi(g)) = g$ . Now, consider some edge  $e = \{g, g'\} \in E^*$  where  $g < g'$ , and denote  $i = i(g)$  and  $j = i(g')$ . Since  $\{g, g'\} \in E^*$ , it holds that the intersection of  $X$  with  $S^e$  cuts  $S^e$  precisely, and therefore  $(i, j) \in I^e$ . By the definition of  $I^e$ , we deduce that  $\{v_i^g, v_j^{g'}\} \in E(G)$ , and thus  $\{\psi(g), \psi(g')\} \in E(H)$ . This concludes the proof of the lemma.  $\square$

As a direct corollary to Lemma 3.15 with  $\epsilon = 0$ , we have the following result.

**Corollary 2.** *If  $(H, G, col)$  is a no-instance of PSI, then  $(D, k, \ell, w)$  is a no-instance of A-DOCT.*

We are now ready to derive the correctness of Theorem 1.

*Proof of Theorem 1.* By Lemma 3.14, Corollary 2 and Observations 3.13 and 3.14, given any instance  $(H, G, col)$  of PSI, we can construct in polynomial time an equivalent instance  $(D, k, \ell, w)$  of A-DOCT such that  $k \leq 121|E(G)|$ . Thus, by Proposition 3.1, A-DOCT is  $W[1]$ -hard. Moreover, by Proposition 3.1, unless ETH fails, A-DOCT cannot be solved in time  $f(k) \cdot n^{o(\frac{k}{\log k})}$  for any function  $f$ . Hence, by Corollary 1, we conclude that Theorem 1 is correct.  $\square$

## 4 Parameterized Inapproximability

In this section, we prove Theorem 3. For convenience, let us restate the theorem below.

**Theorem 3.** *Assuming Gap-ETH or PIH and  $FPT \neq W[1]$ , there exists an  $\epsilon > 0$  such that DOCT does not admit an FPT-approximation algorithm with approximation ratio  $1 + \epsilon$ .*

We first recall basic concepts concerning constraint satisfaction, after which we define the gap-ETH and show that an FPT-approximation algorithm for  $\epsilon$ -gap-BCSP would violate the gap-ETH. Then we reduce  $\epsilon$ -gap-BCSP to the special case of this problem where every variable occurs in at most three constraints. Finally we conclude the proof of Theorem 3.

### 4.1 Constraint Satisfaction

Given a set of variables  $X = \{x_1, x_2, \dots, x_k\}$  and a family of pairwise-disjoint domains  $\mathcal{D} = \{D_1, D_2, \dots, D_k\}$ , a *binary constraint* is a pair  $c = ((x_i, x_j), R)$  where  $x_i, x_j \in X$ ,  $i \neq j$ , and  $R$  is a binary relation over  $D_i \times D_j$ . An *evaluation* is a function  $\psi : X \rightarrow \bigcup \mathcal{D}$  such that for all  $x_i \in X$ ,  $\psi(x_i) \in D_i$ . An evaluation  $\psi$  is said to *satisfy*  $((x_i, x_j), R)$  if  $(\psi(x_i), \psi(x_j)) \in R$ . Moreover, given a multiset  $C$  of binary constraints, an evaluation  $\psi$  is said to *satisfy*  $C$  if it satisfies every constraint  $c \in C$ . For all  $i \in [k]$ , let  $C_i \subseteq C$  denote the sub(multi)set of constraints where  $x_i$  occurs, and let  $s_i = |C_i|$ . We assume w.l.o.g. that for all  $i \in [k]$ ,  $s_i \geq 1$ , that is, for every variable in  $X$ , there exists at least one binary constraint where it occurs.

Here, we cannot assume w.l.o.g. that  $C$  is a set rather than a multiset, or more generally, that for all distinct  $i, j \in [k]$ ,  $|C_i \cap C_j| \leq 1$ .<sup>7</sup> Indeed, here we do not only care whether we can satisfy  $C$  or not (if this were the case, then it would have been trivial to see how, given an instance where  $|C_i \cap C_j| \geq 2$  for some  $i, j \in [k]$ , obtain an equivalent instance with the same set of variables and where  $|C_i \cap C_j| \leq 1$  for all  $i, j \in [k]$ ), but also what is the fraction of constraints that are satisfied. However, as we would see later, it would actually still be simple to ensure the property above. We do not ensure it just yet, as in the reduction in Section 4.2, we implicitly rely on the possibility to allow pairs of variables to occur in more than one constraint.

The BINARY CONSTRAINT SATISFACTION PROBLEM (BCSP) is defined as follows.

BINARY CONSTRAINT SATISFACTION PROBLEM (BCSP)	
<i>Input:</i>	A set $X = \{x_1, x_2, \dots, x_k\}$ of $k$ variables, a family of pairwise-disjoint domains $\mathcal{D} = \{D_1, D_2, \dots, D_k\}$ , and a multiset $C$ of binary constraints.
<i>Question:</i>	Does there exist an evaluation that satisfies $C$ ?

Recall that the promise problem  $\epsilon$ -GAP-BCSP is defined as BCSP where the input instance is promised to either be satisfiable, or have the property that every evaluation satisfies less than  $(1 - \epsilon)$  fraction of the constraints.

### 4.2 Hardness of $\epsilon$ -GAP-BCSP assuming the gap-ETH

First we state the gap-ETH. This definition is taken verbatim from Chalermsook et al. [6]. Before we state the hypothesis, let us recall the definition of 3-SAT. In  $q$ -SAT, we are given a CNF formula  $\phi$  in which each clause consists of at most  $q$  literals, and the goal is to decide whether  $\phi$  is satisfiable. MAX  $q$ -SAT is a maximization version of  $q$ -SAT which asks to compute the maximum number of clauses in  $\phi$  that can be simultaneously satisfied. We will abuse  $q$ -SAT to mean MAX  $q$ -SAT, and for a formula  $\phi$ , we use  $SAT(\phi)$  to denote the maximum number of clauses satisfied by any assignment. The Gap Exponential Time Hypothesis can now be stated in terms of SAT as follows.

<sup>7</sup>In other words, we cannot assume w.l.o.g. that for every pair of variables in  $X$ , there exists at most one binary constraint in  $C$  where both of these variables occur.

**Conjecture 4.1** ((randomized) Gap Exponential-Time Hypothesis (gap-ETH) [20, 39]). *For some constant  $\delta, \rho, c > 0$ , no algorithm can, given a 3-SAT formula  $\phi$  on  $n$  variables and  $m \leq cn$  clauses, distinguish between the following cases correctly with probability  $\geq 2/3$  in  $\mathcal{O}(2^{\delta n})$  time:*

- $SAT(\phi) = m$  and
- $SAT(\phi) < (1 - \rho)m$ .

Next we relate gap-ETH with hardness of  $\epsilon$ -GAP-BCSP, the problem underlying PIH.

**Theorem 4.** *Assuming the gap-ETH there exists an  $\epsilon > 0$  such that there is no FPT-approximation algorithm for  $\epsilon$ -GAP-BCSP.*

*Proof.* Given  $\rho$ , an integer  $k$  and a SAT formula  $\phi$  we construct an instance  $I = (X, \mathcal{D}, C)$  of  $\epsilon$ -GAP-BCSP with  $2k$  variables as follows. We will assume that both the number  $n$  of variables and the number  $m$  of clauses in  $\phi$  is divisible by  $k$ , and that the variables are  $v_0, v_1, \dots, v_{n-1}$  and the clauses are  $Q_0, Q_1, \dots, Q_{m-1}$ . We assume without loss of generality that each clause  $Q_i$  contains precisely 3 literals,  $q_i^1, q_i^2, q_i^3$ .

We group the variables into  $k$  groups of  $n/k$  variables each: for  $0 \leq r \leq k-1$  the  $r$ 'th group consists of  $\{v_{\frac{rn}{k}}, v_{\frac{rn}{k}+1}, \dots, v_{\frac{(r+1)n}{k}-1}\}$ . Similarly we group the clauses into  $k$  groups with  $m/k$  clauses in each group as follows. For  $0 \leq r \leq k-1$  the  $r$ 'th group consists of  $\{Q_{\frac{rm}{k}}, Q_{\frac{rm}{k}+1}, \dots, Q_{\frac{(r+1)m}{k}-1}\}$ .

The set  $X$  of variables of  $I$  is  $\{x_0, \dots, x_{k-1}, y_0, \dots, y_{k-1}\}$  and the domain of each variable in  $x_0, \dots, x_k$  is  $\{0, 1\}^{n/k}$ . The domain of each variable in  $y_0, \dots, y_{k-1}$  is  $\{1, 2, 3\}^{m/k}$ . There is a natural one-to-one correspondence between an assignment of a value in  $\{0, 1\}^{n/k}$  to  $x_r$  and an assignment of truth values to the variables in the  $r$ 'th group of variables  $\phi$ . In particular the  $i$ 'th bit of  $x_r$  is equal to the value of  $v_{\frac{rn}{k}+i}$ . We will interpret an assignment of  $\{1, 2, 3\}^{m/k}$  to a variable  $y_r$  as follows. If the  $i$ 'th character of  $y_r$  is  $t$  then the  $t$ 'th literal of the clause  $Q_{\frac{rm}{k}+i}$  (namely  $q_{\frac{rm}{k}+i}^t$ ) is set to true.

For each clause  $Q_i$  of  $\phi$  and literal  $q_i^t \in Q_i$  we add a constraint to the instance  $I$ . (We remark that it would thus be possible that for a pair of variables, more than one constraint where they occur may be added.) Let  $i = \frac{rm}{k} + j$ , in other words  $Q_i$  is the  $j$ 'th clause of the  $r$ 'th group. Similarly, let  $q_i^t$  be a literal corresponding to the variable  $v_{\frac{sn}{k}+l}$ , either positively or negatively. Thus  $q_i^t$  is a literal of the  $\ell$ 'th variable in the  $s$ 'th variable group of  $\phi$ .

We add a constraint to  $I$  between  $x_s$  and  $y_r$ . If  $q_i^t$  is a positive literal, this constraint is satisfied unless the  $j$ 'th character of  $y_r$  is  $t$  and the  $\ell$ 'th bit of  $x_r$  is 0. If  $q_i^t$  is a negative literal, this constraint is satisfied unless the  $j$ 'th character of  $y_r$  is  $t$  and the  $\ell$ 'th bit of  $x_r$  is 1. In other words, the constraint is falsified if  $y_r$  ‘‘claims’’ that the literal  $q_i^t$  is set to true while  $x_s$  claims that the variable corresponding to  $q_i^t$  is set such that  $q_i^t$  is false. The total number of such constraints is exactly  $3m$ . Thus the number of variables in the constructed instance is  $2k$  and the total size of the instance is  $\mathcal{O}(k3^{cn/k} + m)$ .

For completeness, suppose there is an assignment of truth values to  $v_0, \dots, v_{n-1}$  that satisfies all clauses. Set  $x_i$ 's according to this assignment. For every clause  $Q_i$  there is a literal  $q_i^{t_i}$  that is set to true. Let  $i = \frac{mr}{k} + j$ , set the  $j$ 'th character of  $y_r$  to  $t_i$ . By construction all  $3m$  constraints of the instance  $I$  are satisfied.

For soundness, suppose that there is an assignment to  $x_0, \dots, x_{k-1}$  and  $y_0, \dots, y_{k-1}$  that satisfies at least  $(1 - \frac{\epsilon}{3})3m$  constraints of  $I$ . We consider the assignment to  $v_0, \dots, v_{n-1}$  that corresponds to the assignment to  $x_0, \dots, x_{k-1}$ . We argue that at least  $(1 - \epsilon)m$  clauses are satisfied by this assignment.

There are at most  $\epsilon m$  unsatisfied constraints in  $I$ . Each clause  $Q_i$  of  $\phi$  gave rise to 3 constraints in  $I$ , call the clause  $Q_i$  *happy* if all of its 3 constraints are satisfied. Since there are

at most  $\epsilon m$  unsatisfied constraints, there are at least  $(1 - \epsilon)m$  happy clauses. We now argue that all happy clauses are satisfied.

Let  $Q_i$  be a happy clause, and let  $i = \frac{rm}{k} + j$ . Let  $t_i$  be the  $j$ 'th character of  $y_r$ . Let  $v_{\frac{sn}{k} + \ell}$  be the variable that corresponds to the literal  $q_i^{t_i}$ . In the construction of the instance  $I$  we added a constraint that would be falsified if the  $j$ 'th character of  $y_r$  is  $t_i$  and at the same time the  $\ell$ 'th bit of  $x_s$  was set such that setting  $v_{\frac{sn}{k} + \ell}$  to the  $\ell$ 'th bit of  $x_s$  falsifies the literal  $q_i^{t_i}$ . Since the clause  $Q_i$  is happy this constraint is not falsified and hence the clause  $Q_i$  is satisfied.

Suppose now that  $\epsilon/3$ -GAP-BCSP has an FPT algorithm with running time  $f(k)|I|^{\mathcal{O}(1)}$  where  $|I|$  is the size of the input instance  $I$ . Let  $g(k)$  be a non-decreasing function of  $k$  that tends to infinity with  $k$  such that  $f(2 \cdot g(k)) \leq k$ .

Construct from  $\phi$  a  $\epsilon/3$ -GAP-BCSP instance as defined above with  $k = g(n)$ , run the algorithm for  $\epsilon/3$ -GAP-BCSP, and return the same answer. The total running time of the algorithm is upper bounded by

$$f(2k)|I|^{\mathcal{O}(1)} \leq f(g(n)) \cdot (k3^{cn/k} + m)^{\mathcal{O}(1)} \leq n \cdot g(n) \cdot 3^{\mathcal{O}(cn/g(n))} \cdot m^{\mathcal{O}(1)} \cdot 2^{\mathcal{O}(n)}.$$

This contradicts the gap-ETH, concluding the proof.  $\square$

### 4.3 Constraint Satisfaction with Bounded Degree

For all  $i, j \in [k]$ , denote  $C_{\{i,j\}} = C_i \cap C_j$ , that is, the multiset of constraints in  $C$  where  $x_i$  and  $x_j$  occur together (if a constraint  $c$  belongs to  $C_{\{i,j\}}$ , then its number of occurrences is equal to its number of occurrences in  $C$ ). Moreover, for all  $i, j \in [k]$ , denote  $s_{\{i,j\}} = |C_{\{i,j\}}|$ . Note that  $|C| = \sum_{i,j \in [k], i < j} s_{\{i,j\}}$ . For a fixed integer  $d$ , the  $\epsilon$ -GAP-BCSP $_d$  is defined as the special case of  $\epsilon$ -GAP-BCSP where every variable is present in at most  $d$  constraints, and for all  $i, j \in [k]$ ,  $s_{\{i,j\}} \leq 1$  (that is, every two variables occur together in at most one constraint). Before we turn to prove the hardness of  $\epsilon$ -GAP-BCSP $_3$ , which is the main part of this section, we first simplify the instances of  $\epsilon$ -GAP-BCSP that we need to analyze.

We prove that in what follows, it would be ‘‘safe’’ to assume that  $|C| = k^{\mathcal{O}(1)}$ . To this end, we first prove the following simple lemma.

**Lemma 4.1.** *There exists a polynomial-time algorithm that, given an instance  $I = (X, \mathcal{D}, C)$  of  $\epsilon$ -GAP-BCSP with  $k = |X|$ , constructs another instance  $\hat{I} = (X, \mathcal{D}, \hat{C})$  of  $\epsilon$ -GAP-BCSP with  $|\hat{C}| = |C|$  and the following properties:*

- for all  $i, j \in [k]$ , there is at most one distinct constraint in  $\hat{C}_{\{i,j\}}$ ;<sup>8</sup>
- if  $I$  is satisfiable, then  $\hat{I}$  is satisfiable;
- if  $\hat{I}$  admits an evaluation that satisfies at least  $(1 - \epsilon)|\hat{C}|$  constraints, then  $I$  admits an evaluation that satisfies at least  $(1 - \epsilon)|C|$  constraints.

*Proof.* Let  $I = (X, \mathcal{D}, C)$  be an instance of  $\epsilon$ -GAP-BCSP with  $k = |X|$ . For all  $i, j \in [k]$ ,  $i < j$ , denote  $\hat{R}_{i,j} = \bigcap_{c=((x_i, x_j), R) \in C} R$ ,  $\hat{R}_{j,i} = \bigcap_{c=((x_j, x_i), R) \in C} R$  and  $\hat{R}_{\{i,j\}} = \hat{R}_{i,j} \cap \{(d, d') : (d', d) \in R_{j,i}\}$ . We define  $\hat{C}$  as follows. For all  $i, j \in [k]$ ,  $i < j$ , the multiset  $\hat{C}$  contains the constraint  $c_{\{i,j\}} = ((x_i, x_j), R_{\{i,j\}})$  exactly  $s_{\{i,j\}}$  times. Clearly,  $|\hat{C}| = |C|$ , and for all  $i, j \in [k]$ , there is at most one distinct constraint in  $\hat{C}_{\{i,j\}}$ . Moreover, every evaluation of  $I$  that satisfies  $C$  also satisfies  $\hat{C}$ . As every evaluation that satisfies the distinct constraint in  $\hat{C}_{\{i,j\}}$  also satisfies all constraints  $C_{\{i,j\}}$ , we further have that if  $\hat{I}$  admits an evaluation that satisfies at least  $(1 - \epsilon)|\hat{C}|$  constraints, then  $I$  admits an evaluation that satisfies at least  $(1 - \epsilon)|C|$  constraints.  $\square$

<sup>8</sup>The distinct constraint in  $\hat{C}_{\{i,j\}}$  may occur more than once in  $\hat{C}_{\{i,j\}}$ , that is,  $\hat{s}_{\{i,j\}}$  may be larger than 1.

We now show how to ensure that  $|C|$  is small.

**Lemma 4.2.** *There exists a polynomial-time algorithm that, given an instance  $I = (X, \mathcal{D}, C)$  of  $\epsilon$ -GAP-BCSP with  $k = |X|$ , constructs another instance  $\hat{I} = (X, \mathcal{D}, \hat{C})$  of  $\frac{\epsilon}{2}$ -GAP-BCSP with the following properties:*

- $|\hat{C}| \leq \frac{k^4}{\epsilon}$ ;
- if  $I$  is satisfiable, then  $\hat{I}$  is satisfiable;
- if  $\hat{I}$  admits an evaluation that satisfies at least  $(1 - \frac{\epsilon}{2})|\hat{C}|$  constraints, then  $I$  admits an evaluation that satisfies at least  $(1 - \epsilon)|C|$  constraints.

*Proof.* Let  $I' = (X, \mathcal{D}, C')$  be an instance of  $\epsilon$ -GAP-BCSP with  $k = |X|$ . Suppose that  $|C| > \frac{k^4}{\epsilon}$ , else we are already done. By Lemma 4.1, we obtain in polynomial time an instance  $I = (X, \mathcal{D}, C)$  such that for all  $i, j \in [k]$ , there is at most one distinct constraint in  $C_{\{i,j\}}$ , and  $|C| = |C'|$ . Now, we construct in polynomial time an instance  $\hat{I} = (X, \mathcal{D}, \hat{C})$  of  $\frac{\epsilon}{2}$ -GAP-BCSP as follows. Denote  $K = \binom{k}{2}$  and  $T = \frac{\epsilon}{2}|C| (> \frac{k^4}{2})$ . Let  $\hat{C}_{\{i,j\}}$  be a multiset of  $\hat{s}_{\{i,j\}} := \lfloor \frac{K}{T} s_{\{i,j\}} \rfloor (\leq s_{\{i,j\}})$  occurrences of the distinct constraint in  $C_{\{i,j\}}$ . Define  $\hat{C} = \bigcup_{i,j \in [k], i < j} \hat{C}_{\{i,j\}}$  (note that  $\hat{C}$  is a multiset).

Observe that  $|\hat{C}| = \sum_{i,j \in [k], i < j} \hat{s}_{\{i,j\}} = \sum_{i,j \in [k], i < j} \lfloor \frac{K}{T} s_{\{i,j\}} \rfloor \leq \sum_{i,j \in [k], i < j} \frac{K}{\frac{\epsilon}{2}|C|} s_{\{i,j\}}$ . Since for all  $i, j \in [k]$ ,  $s_{\{i,j\}} \leq |C|$ , we have that  $|\hat{C}| \leq \sum_{i,j \in [k], i < j} \frac{K}{\frac{\epsilon}{2}} = \frac{2K^2}{\epsilon} \leq \frac{k^4}{\epsilon}$  as required. Moreover, because  $\hat{C} \subseteq C$  and by Lemma 4.1, we have that if  $I'$  is satisfiable, then  $I$  is satisfiable, and therefore  $\hat{I}$  is satisfiable as well.

Now, suppose that  $\hat{I}$  admits an evaluation  $\psi$  that satisfies at least  $(1 - \frac{\epsilon}{2})|\hat{C}|$  constraints. By Lemma 4.1, to show that  $I'$  admits an evaluation that satisfies at least  $(1 - \epsilon)|C'|$  constraints, it suffices to show that  $I$  admits an evaluation that satisfies at least  $(1 - \epsilon)|C|$  constraints. Let  $A$  be the collection of sets  $\{i, j\}$ ,  $i, j \in [k]$  where  $i < j$ , such that  $\psi$  satisfies the distinct constraint in  $C_{\{i,j\}}$ . Then, the number of constraints in  $C$  that  $\psi$  satisfies is  $\sum_{\{i,j\} \in A} s_{\{i,j\}} \geq \frac{T}{K} \sum_{\{i,j\} \in A} \hat{s}_{\{i,j\}} \geq \frac{T}{K} (1 - \frac{\epsilon}{2})|\hat{C}| \geq \frac{T}{K} (1 - \frac{\epsilon}{2}) \sum_{i,j \in [k], i < j} \hat{s}_{\{i,j\}} \geq \frac{T}{K} (1 - \frac{\epsilon}{2}) \sum_{i,j \in [k], i < j} (\frac{K}{T} s_{\{i,j\}} - 1) = \frac{T}{K} (1 - \frac{\epsilon}{2}) (\frac{K}{T} |C| - K) = (1 - \frac{\epsilon}{2})|C| - (1 - \frac{\epsilon}{2})T \leq (1 - \frac{\epsilon}{2})|C| - T = (1 - \frac{\epsilon}{2})|C| - \frac{\epsilon}{2}|C| = (1 - \epsilon)|C|$ . This concludes the proof.  $\square$

In light of Lemma 4.2, throughout the remainder of this section we assume that  $|C| \leq \frac{k^4}{\epsilon}$ .

Recall that for a fixed integer  $d$ , the  $\epsilon$ -GAP-BCSP $_d$  is defined as the special case of  $\epsilon$ -GAP-BCSP where every variable is present in at most  $d$  constraints, and for all  $i, j \in [k]$ ,  $s_{\{i,j\}} \leq 1$ . Towards the proof of the hardness of  $\epsilon$ -GAP-BCSP $_3$ , we first consider  $\epsilon$ -GAP-BCSP $_4$ . For this proof, we need to recall the notion of an expander.

**Definition 4.1.** *Given  $n, d \in \mathbb{N}$  and  $0 \leq \gamma \leq 1$ , an  $(n, d, \gamma)$ -expander is an undirected  $d$ -regular graph  $G$  on  $n$  vertices such that for every set  $S \subseteq V(G)$  of size at most  $\frac{1}{2}|V(G)|$ , the number of edges with one endpoint in  $S$  and the other endpoint in  $V(G) \setminus S$  is at least  $\gamma \cdot d \cdot |S|$ .*

For the sake of brevity, given  $n_1, n_2 \in \mathbb{N}$ , we refer to any  $(n, d, \gamma)$ -expander where  $n_1 \leq n \leq n_2$  as an  $([n_1, n_2], d, \gamma)$ -expander. We would also need to rely on the following result.

**Proposition 4.1** ([2]). *There exist  $\gamma > 0$  and  $\ell \in \mathbb{N}$  such that for all  $s \in \mathbb{N}$ , an  $([s, \ell s], 3, \gamma)$ -expander can be constructed in polynomial time.*

**Lemma 4.3.** *Assuming PIH, there exists an  $\epsilon > 0$  such that  $\epsilon$ -GAP-BCSP $_4$  is W[1]-hard.*

*Proof.* To prove that the lemma is correct, we present a reduction from  $\epsilon$ -GAP-BCSP to  $\delta$ -GAP-BCSP<sub>4</sub> where  $\delta = \frac{\epsilon}{4\ell(1 + \frac{1}{3\gamma})}$ . Here,  $\gamma$  and  $\ell$  are the fixed constants stated in Proposition

4.1. For this purpose, let  $I = (X, \mathcal{D}, C)$  be an instance of  $\epsilon$ -GAP-BCSP. Then, we construct an instance  $\widehat{I} = (\widehat{X}, \widehat{\mathcal{D}}, \widehat{C})$  of  $\delta$ -GAP-BCSP<sub>4</sub> as follows. First, for all  $i \in [k]$ , apply Proposition 4.1 to construct an  $([s_i, \ell s_i], 3, \gamma)$ -expander  $G_i$  (recall that  $s_i = |C_i|$ , the number of constraints where  $x_i$  occurs), and denote  $n_i = |V(G_i)|$ . Now, for all  $i \in [k]$ , define  $\widehat{X}^i = \{\widehat{x}_1^i, \widehat{x}_2^i, \dots, \widehat{x}_{n_i}^i\}$ , and let  $f^i : \widehat{X}^i \rightarrow V(G_i)$  be an arbitrarily chosen bijective function. Accordingly, set  $\widehat{X} = \bigcup_{i=1}^k \widehat{X}^i$ . Recall that we assume that  $|C| \leq \frac{k^4}{\epsilon}$ . Moreover, since the constraints are binary, we have that  $\sum_{i=1}^k s_i = 2|C|$ . Therefore,  $|\widehat{X}| = \sum_{i=1}^k n_i \leq \ell \sum_{i=1}^k s_i = 2\ell|C| \leq 2\ell \frac{k^4}{\epsilon}$ . That is, the number of variables in the new instance is bounded by a function of  $k$ .

We proceed to define  $\widehat{\mathcal{D}}$  by letting the domain of every  $\widehat{x}_j^i$ , where  $i \in [k]$  and  $j \in [n_i]$ , be  $\widehat{D}_j^i = \{\widehat{d}_j^i : d \in D_i\}$ . Finally, let us define  $\widehat{C}$  as follows. For all  $i \in [k]$ , let  $g^i : C_i \rightarrow X^i$  be an arbitrarily chosen injective function. Then, for all  $c = ((x_i, x_j), R) \in C$ , denote  $\widehat{c} = ((\widehat{x}_p^i, \widehat{x}_q^i), \widehat{R}) = \{(\widehat{d}_p^i, \widehat{d}_q^i) : (d, d') \in R\}$  where  $\widehat{x}_p^i = g^i(x_i)$  and  $\widehat{x}_q^i = g^i(x_j)$ . Define  $C^* = \{\widehat{c} : c \in C\}$ . We now define a set of equality constraints on the set of ‘copies’ of each variable. Define  $C_- = \{((\widehat{x}_p^i, \widehat{x}_q^i), \{(d_p^i, d_q^i) : d \in D_i\}) : i \in [k], p, q \in [n_i], (p, q) \in E(G_i)\}$ . Finally, set  $\widehat{C} = C^* \cup C_-$ . Since for all  $i \in [k]$ , the graph  $G_i$  is 3-regular, we have that every variable in  $\widehat{X}^i$  occurs in at most three constraints in  $C_-$ . Moreover, since for all  $i \in [k]$  the function  $g^i$  is injective, we have that every variable in  $\widehat{X}^i$  occurs in at most one constraint in  $C^*$ . Thus, every variable in  $\widehat{X}$  occurs in at most four constraints in total. Thus, since every constraint is binary, we also have that  $|\widehat{C}| \leq 2|\widehat{X}|$ . Since we have already shown that  $|\widehat{X}| \leq 2\ell|C|$ , we derive that  $|\widehat{C}| \leq 4\ell|C|$ .

To prove that the reduction is correct, we argue that if  $I$  is satisfiable (admits an evaluation satisfying all constraints) then so is  $\widehat{I}$  and conversely, if  $\widehat{I}$  admits an evaluation that satisfies at least  $(1 - \delta)|\widehat{C}|$  constraints, then  $I$  admits an evaluation that satisfies at least  $(1 - \epsilon)|C|$  constraints.

Suppose that  $I$  admits an evaluation  $\psi$  that satisfies all of the constraints. Then, we have that  $\widehat{I}$  also admits an evaluation  $\widehat{\psi}$  that satisfies all of the constraints. Indeed, we simply define  $\widehat{\psi}$  by setting  $\widehat{\psi}(\widehat{x}_j^i) = d_j^i$ , where  $d = \psi(x_i)$ , for all  $i \in [k]$  and  $j \in [n_i]$ .

Next, suppose that  $\widehat{I}$  admits an evaluation  $\widehat{\psi}$  that satisfies at least  $(1 - \delta)|\widehat{C}|$  constraints. We define an evaluation  $\psi$  for  $I$  as follows. For all  $i \in [k]$  and  $d \in D_i$ , define  $X^i(d) = \{\widehat{x}_j^i \in X^i : \widehat{\psi}(\widehat{x}_j^i) = d_j^i\}$ . Now, for all  $i \in [k]$ , let  $\widetilde{d}^i$  be a value in  $D_i$  that among all values in  $D_i$ , maximizes  $|X^i(d)|$  (if there is more than one choice, choose one arbitrarily). Moreover, for all  $i \in [k]$ , denote  $Y^i = X^i \setminus X^i(\widetilde{d}^i)$ . Then, for all  $i \in [k]$ , we set  $\psi(x_i) = \widetilde{d}^i$ . We claim that  $\psi$  satisfies at least  $(1 - \epsilon)|C|$  of the constraints in  $C$ . To show this, we first note that all  $\widehat{c} = ((\widehat{x}_p^i, \widehat{x}_q^i), R)$  such that either both  $\widehat{\psi}(\widehat{x}_p^i) = \widetilde{d}^i$  and  $\widehat{\psi}(\widehat{x}_q^i) \neq \widetilde{d}^i$  or both  $\widehat{\psi}(\widehat{x}_p^i) \neq \widetilde{d}^i$  and  $\widehat{\psi}(\widehat{x}_q^i) = \widetilde{d}^i$ , a unique constraint in  $C_-$  is violated by  $\widehat{\psi}$ . Since for all  $i \in [k]$ ,  $G_i$  is an  $([s_i, \ell s_i], 3, \gamma)$ -expander, we

derive that at least  $\sum_{i=1}^k \gamma \cdot 3 \cdot |Y^i| = 3\gamma \sum_{i=1}^k |Y^i|$  constraints in  $C_-$  are violated by  $\widehat{\psi}$ . Thus,

$3\gamma \sum_{i=1}^k |Y^i| < \delta|\widehat{C}|$ , which implies that  $\sum_{i=1}^k |Y^i| \leq \frac{\delta}{3\gamma}|\widehat{C}|$ . Let us now denote by  $C_Y$  the set of all

constraints  $c = ((x_i, x_j), R) \in C$  such that  $g^i(c) \in Y^i$  or  $g^j(c) \in Y_j$ . Then, since for all  $i \in [k]$ ,  $g^i$  is an injective function, we have that  $|C_Y| \leq \sum_{i=1}^k |Y^i|$ , and thus we derive that  $|C_Y| \leq \frac{\delta}{3\gamma}|\widehat{C}|$ .

Notice that for all  $c \in C \setminus C_Y$  that is violated by  $\psi$ , there is a unique constraint in  $C^*$  that is also violated by  $\hat{\psi}$ . Thus, the number of constraints in  $C \setminus C_Y$  that are violated by  $\psi$  is smaller than  $\delta|\hat{C}|$ . Overall, we conclude that  $\psi$  violates less than  $\delta|\hat{C}| + |C_Y| \leq (1 + \frac{1}{3\gamma})\delta|\hat{C}| \leq 4\ell(1 + \frac{1}{3\gamma})\delta|C|$  constraints in  $C$ . Since  $\delta = \frac{\epsilon}{4\ell(1 + \frac{1}{3\gamma})}$ , we conclude that  $\psi$  violates less than  $\epsilon|C|$  constraints in  $C$ . This concludes the proof of the lemma.  $\square$

We now turn to prove the hardness of  $\epsilon$ -GAP-BCSP<sub>3</sub>.

**Lemma 4.4.** *Assuming PIH, there exists an  $\epsilon > 0$  such that  $\epsilon$ -GAP-BCSP<sub>3</sub> is W[1]-hard.*

*Proof.* By Lemma 4.3, we have that there exists an  $\epsilon > 0$  such that  $\epsilon$ -GAP-BCSP<sub>4</sub> is W[1]-hard. To prove that the lemma is correct, we present a reduction from  $\epsilon$ -GAP-BCSP<sub>4</sub> to  $\delta$ -GAP-BCSP<sub>3</sub> where  $\delta = \epsilon/15$ . For this purpose, let  $I = (X, \mathcal{D}, C)$  be an instance of  $\epsilon$ -GAP-BCSP<sub>4</sub>. Then, we construct an instance  $\hat{I} = (\hat{X}, \hat{\mathcal{D}}, \hat{C})$  of  $\delta$ -GAP-BCSP<sub>3</sub> as follows. First, define  $\hat{X} = X \cup X'$  where  $X' = \{x'_1, x'_2, \dots, x'_k\}$ . Hence,  $|\hat{X}| \leq 2k$ . Now, let us define  $\hat{\mathcal{D}}$ . For all  $i \in [k]$ , the domain of  $x_i$  is defined as  $D_i \in \mathcal{D}$ , and the domain of  $x'_i$  is defined as  $D'_i = \{d' : d \in D_i\}$ . Now, for all  $i \in [k]$ , let  $(A_i, B_i)$  be a partition of  $C_i$  such that  $|A_i|, |B_i| \leq 2$ . Note that the existence of such a partition follows from the fact that  $|C_i| \leq 4$ . Then, for all  $c = ((x_i, x_j), R) \in C$ , denote  $\hat{c} = ((x_i, x_j), R)$  if  $((x_i, x_j), R) \in A_i \cap A_j$ ,  $\hat{c} = ((x_i, x'_j), R' = \{(p, q') : (p, q) \in R\})$  if  $((x_i, x_j), R) \in A_i \cap B_j$ ,  $\hat{c} = ((x'_i, x_j), R' = \{(p', q) : (p, q) \in R\})$  if  $((x_i, x_j), R) \in B_i \cap A_j$  and  $\hat{c} = ((x'_i, x'_j), R' = \{(p', q') : (p, q) \in R\})$  otherwise. Define  $C^* = \{\hat{c} : c \in C\}$ . We now define a set of equality constraints on each pair of copies of the variables. Define  $C_- = \{((x_i, x'_i), \{(d, d') : d \in D_i\}) : i \in [k]\}$ . Finally, set  $\hat{C} = C^* \cup C_-$ . Clearly, every variable in  $\hat{X}$  occurs in at most three constraints in  $\hat{C}$ , and the construction can be performed in polynomial time.

To prove that the reduction is correct, we argue that if  $I$  is satisfiable then so is  $\hat{I}$  and conversely, if  $\hat{I}$  admits an evaluation that satisfies at least  $(1 - \delta)|\hat{C}|$  constraints, then  $I$  admits an evaluation that satisfies at least  $(1 - \epsilon)|C|$  constraints.

Suppose that  $I$  admits an evaluation  $\psi$  that satisfies all of the constraints. Then, we have that  $\hat{I}$  also admits an evaluation  $\hat{\psi}$  that satisfies all of the constraints. Indeed, we simply define  $\hat{\psi}$  by setting  $\hat{\psi}(x_i) = \psi(x_i)$  and  $\hat{\psi}(x'_i) = \psi(x_i)'$  for all  $i \in [k]$ .

Next, suppose that  $\hat{I}$  admits an evaluation  $\hat{\psi}$  that satisfies at least  $(1 - \delta)|\hat{C}|$  constraints. We define an evaluation  $\psi$  for  $I$  as follows. For all  $i \in [k]$ , we set  $\psi(x_i) = \hat{\psi}(x_i)$ . We claim that  $\psi$  satisfies at least  $(1 - \epsilon)|C|$  of the constraints in  $C$ . To show this, we denote  $Y = \{x_i \in X : \hat{\psi}(x_i) \neq \hat{\psi}(x'_i)\}$ . Since  $\hat{\psi}$  violates less than  $\delta|\hat{C}|$  constraints, we have that  $|Y| < \delta|\hat{C}|$ . Let  $C_Y$  denote the subset of constraints of  $C$  where at least one variable of  $Y$  occurs. Note that since every variable in  $X$  occurs in at most four constraints in  $C$ , we have that  $|C_Y| \leq 4|Y| \leq 4\delta|\hat{C}|$ . Moreover, note that for every constraint  $\hat{c} \in C^* \setminus C_Y$  that is satisfied by  $\hat{\psi}$  is also satisfied by  $\psi$ . Since  $\hat{\psi}$  violates less than  $\delta|\hat{C}|$  constraints in total, we have that  $\hat{\psi}$  also violates less than  $\delta|\hat{C}|$  constraints from  $C^* \setminus C_Y$ . Thus, we have that  $\psi$  violates less than  $5\delta|\hat{C}|$  constraints from  $C$ . Since every variable occurs in at least one constraint, we have that  $|\hat{C}| = |C^*| + |C_-| = |C| + |X| \leq 3|C|$ . We thus conclude that  $\psi$  violates less than  $15\delta|C| = \epsilon|C|$  constraints from  $C$ . This concludes the proof of the lemma.  $\square$

#### 4.4 Proof of Theorem 3

We now translate Hypothesis 1 in terms of PSI. For this purpose, we define the promise problem  $\epsilon$ -GAP-PSI as PSI where the input instance is promised to either be a yes-instance, or have the property that for every subgraph  $G'$  of  $G$  with at least  $(1 - \epsilon)|E(G)|$  edges, there does not exist a colorful mapping of  $G'$  into  $H$ . It is straightforward to see that if  $\epsilon$ -GAP-BCSP is W[1]-hard, then  $\epsilon$ -GAP-PSI is W[1]-hard as well. For the sake of completeness, we present the reduction.

**Lemma 4.5.** *Assuming PIH, there exists an  $\epsilon > 0$  such that  $\epsilon$ -GAP-PSI is W[1]-hard.*

*Sketch.* Let  $(X, \mathcal{D}, C)$  be an instance of  $\epsilon$ -GAP-BCSP<sub>3</sub>. Then, we construct (in polynomial time) an instance  $(H, G, col)$  of PSI as follows. First, we set  $V(G) = X$  and  $E(G) = \{\{x, x'\} : \text{there exists } R \text{ such that } ((x, x'), R) \in C\}$ . Second, we set  $V(H) = \bigcup \mathcal{D}$ , and we let  $E(H)$  contain every edge  $\{d, d'\}$  for which there exist  $i \neq j$  and  $R$  such that  $d \in D_i, d' \in D_j, ((x_i, x_j), R) \in C$  and  $(d, d') \in R$ . Finally, for all  $i \in [k]$  and  $d \in D_i$ , we set  $col(d) = x_i$ . Notice that for all  $0 \leq \alpha \leq 1$ , there exists  $C' \subseteq C$  of size at least  $\alpha|C|$  such that there exists an evaluation satisfying  $C'$  if and only if there exists a subgraph  $G'$  of  $G$  with at least  $\alpha|E(G)|$  edges such that there exists a colorful mapping of  $G'$  into  $H$ . Hence, by Lemma 4.4, we conclude that the lemma is correct.  $\square$

We remark that it is also straightforward to see that if  $\epsilon$ -GAP-PSI is W[1]-hard, then  $\epsilon$ -GAP-BCSP<sub>3</sub> is W[1]-hard, and hence  $\epsilon$ -GAP-BCSP is W[1]-hard as well.

Finally, we are ready to prove the correctness of Theorem 3.

*Proof of Theorem 3.* Suppose that there exists an  $\epsilon > 0$  for which there does not exist an FPT algorithm for  $\epsilon$ -GAP-PSI. Let  $\delta = \delta(\epsilon)$  be defined according to Lemma 3.15. We claim that A-DOCT does not admit a  $(1 + \delta)$ -approximation algorithm that runs in time  $f(k) \cdot n^{\mathcal{O}(1)}$  for any function  $f$ . By Lemma 4.5 and Corollary 1, we would thus conclude the correctness of Theorem 3. Suppose, by way of contradiction, that our claim is false. Then, let  $\mathcal{B}$  be a  $(1 + \delta)$ -approximation algorithm for A-DOCT that runs in time  $f(k) \cdot n^{\mathcal{O}(1)}$  for some function  $f$ . We define an algorithm, Algorithm  $\mathcal{A}$  as follows. Given an instance  $(H, G, col)$  of PSI, it constructs the instance  $(D, k, \ell, w)$  of A-DOCT as described in Section 3.5, and calls Algorithm  $\mathcal{B}$  with  $(D, k, \ell, w)$  as input. Then, if  $\mathcal{B}$  outputs NO, then  $\mathcal{A}$  outputs NO, and otherwise  $\mathcal{A}$  outputs YES. By Observations 3.13 and 3.14, Algorithm  $\mathcal{A}$  runs in time  $f(|E(H)|) \cdot |I|^{\mathcal{O}(1)}$  where  $|I|$  is the size of the input instance. On the one hand, by Lemma 3.14, if Algorithm  $\mathcal{A}$  is given as input a yes-instance of PSI, then it constructs a yes-instance of A-DOCT. Next, since  $\mathcal{B}$  is a  $(1 + \delta)$ -approximation algorithm for A-DOCT, Algorithm  $\mathcal{A}$  outputs YES. On the other hand, suppose that Algorithm  $\mathcal{A}$  is given as input an instance  $(H, G, col)$  of PSI for which there does not exist a subgraph  $G'$  of  $G$  with at least  $(1 - \epsilon)|E(G)|$  edges such that there exists a colorful mapping of  $G'$  into  $H$ . By Lemma 3.15, Algorithm  $\mathcal{A}$  constructs an instance  $(D, k, \ell, w)$  of A-DOCT which does not admit a  $(1 + \delta)$ -approximate solution. Then, since  $\mathcal{B}$  is a  $(1 + \delta)$ -approximation algorithm for A-DOCT, Algorithm  $\mathcal{A}$  outputs NO. We have thus reached a contradiction to the choice of  $\epsilon$ . This concludes the proof of Theorem 3.  $\square$

## 5 Parameterized Approximation

We are now ready to begin the section on the parameterized approximation algorithm for DOCT. We will in fact prove a more general result by giving a parameterized approximation algorithm for a covering problem on labeled digraphs that generalizes both DOCT and the NODE UNIQUE LABEL COVER problem [10]. Before we formally define labeled digraphs, we need the following notation.

For  $\ell \in \mathbb{N}$ ,  $S_\ell$  denotes the symmetric group, which is the set of all permutations of the set  $[\ell]$ . For  $\sigma_1, \sigma_2 \in S_\ell$ , we denote by  $\sigma_1 \circ \sigma_2$  the permutation obtained by composing  $\sigma_1$  and  $\sigma_2$  as follows. For every  $i \in [\ell]$ ,  $\sigma_1 \circ \sigma_2(i) = \sigma_1(\sigma_2(i))$ . Finally, unless otherwise specified, all paths and walks we refer to in this section are directed.

Let  $D$  be a digraph with possible self-loops. A *cycle cover* of  $D$  is a set  $\{C_1, \dots, C_r\}$  of vertex-disjoint cycles such that every vertex is part of some cycle. That is,  $V(D) = \bigcup_{i \in [r]} V(C_i)$ . We point out that cycles of length 1 that correspond to self-loops are allowed to be part of a cycle

cover. It is easy to see that the cycle covers of  $V(D)$  are in one-to-one correspondence with the permutations of  $V(D)$ .

**Definition 5.1.** Let  $D$  be a digraph and  $\sigma : A(D) \rightarrow S_\ell$  be an assignment of permutations of  $[\ell]$  to the arcs of  $D$ . We call the pair  $(D, \sigma)$  a labeled digraph. When  $\sigma$  is clear from the context, we just use  $D$  to denote the labeled digraph  $(D, \sigma)$ . For a set  $Z \subseteq V(D)$ , we denote by  $\sigma|_Z$ , the restriction of  $\sigma$  to arcs with both endpoints in  $Z$ . For a directed walk  $P = v_1, \dots, v_r$  in  $D$ , we denote by  $\sigma(P)$  the composed permutation  $\sigma((v_1, v_2)) \circ \dots \circ \sigma((v_{r-1}, v_r))$ .

In all labeled digraphs we work with in this section, there are no *duplicate* arcs. That is, there is no pair  $a, a' = (u, v) \in A(D)$  such that  $\sigma(a) = \sigma(a')$ . However, we may have multiple arcs from  $u$  to  $v$  labeled with distinct elements of  $S_\ell$ . We are now ready to define our main combinatorial structure that, as we will show, generalizes odd cycles in directed graphs.

**Definition 5.2.** Let  $(D, \sigma)$  be a labeled digraph. Let  $H$  be a strongly connected subgraph of  $D$  and let  $v \in V(H)$ . We say that  $H$  is a  $v$ -colorful walk if for every  $i \in [\ell]$ ,  $H$  contains a closed  $v$ -walk  $W$  such that  $\sigma(W)(i) \neq i$ .

**Definition 5.3.** Let  $(D, \sigma)$  be a labeled digraph and  $H$  be a strongly connected subgraph of  $D$ . We say that  $H$  is a colorful walk if it is a  $v$ -colorful walk for every  $v \in V(H)$ . A set  $S \subseteq V(D)$  intersecting every colorful walk contained in  $D$  is called a colorful walk cover of  $D$ . The explicit reference to  $D$  is ignored when  $D$  is clear from the context.

Before we proceed, we provide a short proof of the fact that colorful walks generalize odd cycles and odd closed walks in digraphs. Let  $\ell = 2$  and let  $\pi$  denote the permutation  $(1\ 2)$ . That is,  $\pi(1) = 2$  and  $\pi(2) = 1$ . Furthermore, suppose that for a digraph  $D$ , we assign  $\sigma(a) = \pi$  for every  $a \in A(D)$ . Then, observe that  $D$  has a colorful walk if and only if it has an odd closed walk. Therefore, colorful walks generalize odd closed walks and since a directed graph has an odd cycle if and only if it has an odd closed walk, the following problem is a generalization of DOCT.

COLORFUL WALK COVER

Parameter:  $k, \ell$ .

**Input:** A digraph  $D$ , function  $\sigma : A(D) \rightarrow S_\ell$ , and integer  $k$ .

**Question:** Does  $D$  have a colorful walk cover of size at most  $k$ ?

Due to Theorem 1 and the fact that DOCT is a special case of COLORFUL WALK COVER when  $\ell = 2$ , we get the following corollary.

**Corollary 3.** The COLORFUL WALK COVER problem is  $W[1]$ -hard even for  $\ell = 2$ . Furthermore, assuming the ETH there is no algorithm for COLORFUL WALK COVER with running time  $f(k, \ell)n^{o(k/\log k) \cdot g(\ell)}$  for any functions  $f$  and  $g$ .

We now complement this negative result with a positive approximation result. Recall that due to [31], we know that there is no constant factor approximation for the *optimization* version of DOCT and hence for COLORFUL WALK COVER assuming the Unique Games Conjecture. However, we show that if allowed FPT time, then the optimization version of even COLORFUL WALK COVER can be approximated up to a constant factor. Recall that an algorithm for the optimization version of COLORFUL WALK COVER is an  $\alpha$ -approximation algorithm if it always outputs either a colorful walk cover of size at most  $\alpha \cdot k$  or NO, where if the input instance is a yes-instance, then the algorithm cannot output NO. We now state this result formally.

**Theorem 5.** COLORFUL WALK COVER admits an  $\ell^{\mathcal{O}(k+\ell)} 2^{\mathcal{O}(k^2)} n^{\mathcal{O}(1)}$  time FPT-approximation algorithm with approximation ratio 2.

We obtain our 2-approximation algorithm (Theorem 2) for the optimization version of DOCT as a consequence of this result. The rest of this section is therefore dedicated to proving Theorem 5. The crux of our approximation algorithm is an FPT algorithm (see Lemma 5.1) for the following variant of the COLORFUL WALK COVER problem.

**RESTRICTED COLORFUL WALK COVER (RESTRICTED CWC)**      **Parameter:**  $k, \ell$ .  
**Input:** A digraph  $D$ , function  $\sigma : A(D) \rightarrow S_\ell$ , integer  $k$  and a colorful walk cover  $W \subseteq V(D)$  of size at most  $2k + 1$  such that  $D[W]$  is strongly connected.  
**Question:** Does there exist a colorful walk cover of size at most  $k$  that is disjoint from  $W$ ?

**Lemma 5.1.** RESTRICTED CWC has an algorithm running in time  $\ell^{\mathcal{O}(k+\ell)} 2^{\mathcal{O}(k^2)} n^{\mathcal{O}(1)}$ .

Note that for an instance  $(D, \sigma, k, W)$  of RESTRICTED COLORFUL WALK COVER, we refer to a set  $S$  that is a colorful walk cover of size at most  $k$  and disjoint from  $W$ , as a *solution* for this instance. The rest of the section is organized as follows. In the first subsection, we recall the notions of *shadows* and *shadow covering* and describe the relation between the COLORED WALK COVER problem and the well-studied NODE UNIQUE LABEL COVER (NODE ULC) problem. This relation will play an important role in our algorithm for RESTRICTED COLORFUL WALK COVER. The second subsection is dedicated to a problem-specific *shadow removal* subroutine and the full proof of Lemma 5.1. In the final subsection, we prove Theorem 5 using Lemma 5.1.

### 5.1 Shadow Covering and Connections to NODE UNIQUE LABEL COVER

We begin with the following lemma showing that any subgraph that is a colorful walk with respect to some vertex in it is in fact a colorful walk with respect to every vertex in it.

**Lemma 5.2.** Let  $(D, \sigma)$  be a labeled digraph and  $H$  be a strongly connected subgraph of  $D$ . Then,  $H$  is a colorful walk if and only if it is a  $v$ -colorful walk for some  $v \in V(H)$ .

*Proof.* The forward direction is trivial and hence we now argue the converse direction. That is, suppose that  $H$  is a  $v$ -colorful walk for some  $v \in V(H)$ . Let  $u \in V(H)$  be a vertex distinct from  $v$  and  $i \in [\ell]$ . In order to prove the lemma, it suffices to demonstrate the presence of a closed  $u$ -walk  $Q$  in  $H$  such that  $\sigma(Q)(i) \neq i$ .

Let  $P_1$  and  $P_2$  be arbitrary  $u$ - $v$  and  $v$ - $u$  paths in  $D$  respectively. Let  $i' = \sigma(P_1)(i)$  and let  $W$  be an arbitrary  $v$ -walk in  $H$  such that  $\sigma(W)(i') \neq i'$ . Since  $H$  is a  $v$ -colorful walk,  $W$  exists. Now, consider the two closed  $u$ -walks  $W_1 = P_1 + P_2$  and  $W_2 = P_1 + W + P_2$  that are both contained in  $H$ . Observe that  $\sigma(W_1)(i)$  and  $\sigma(W_2)(i)$  are distinct and at least one of these is distinct from  $i$ . This completes the proof of the lemma.  $\square$

We now define an auxiliary digraph that is a natural directed version of a graph defined by Lokshtanov et al. [37]. We then argue that it is possible to test in polynomial time whether for a given labeled digraph  $(D, \sigma)$ , vertices  $u, v \in V(D)$  and  $\alpha, \beta \in [\ell]$ ,  $D$  has a  $u$ - $v$  walk that ‘maps’  $\alpha$  to  $\beta$ . This subroutine will be used in several places in our algorithm, including in the step where we test whether the input has a colorful walk.

**Definition 5.4.** With every labeled digraph  $(D, \sigma)$ , we associate an auxiliary digraph  $H_{D, \sigma}$  (we ignore the reference to  $\sigma$  when clear from the context), which is defined as follows. The vertex set of  $H_D$  is  $\{v_i | v \in V(D), i \in [\ell]\}$ . The arc set of  $H_D$  is defined as follows. For every arc  $a = (u, v)$  and for every  $i \in [\ell]$ , we have an arc  $(u_i, v_{\sigma(a)(i)})$ . That is, we add an arc from  $u_i$  to  $v_j$  where  $j$  is the image of  $i$  under the permutation  $\sigma(a)$ .

**Lemma 5.3.** Let  $(D, \sigma)$  be a labeled digraph and let  $u$  and  $v$  be (not necessarily distinct) vertices in  $V(D)$ . For every  $i, j \in [\ell]$ , there is a  $u$ - $v$  walk  $W$  in  $D$  such that  $\sigma(W)(i) = j$  if and only if there is a  $u_i$ - $v_j$  path  $P$  in the digraph  $H_D$ .

*Proof.* The lemma follows by a straightforward induction on the length of the walk  $W$  in the forward direction and that of the path  $P$  in the converse direction.  $\square$

Due to the above lemma, it is easy to verify whether a digraph  $D$  has a colorful walk.

**Observation 5.1.** *There is a polynomial time algorithm that, given a labeled digraph  $(D, \sigma)$ , decides whether  $D$  has a colorful walk.*

**Definition 5.5.** *Let  $(D, \sigma)$  be a labeled digraph. We define by  $\text{Doubling}(D)$  the labeled digraph  $(D', \sigma')$  obtained as follows. Initially,  $D' = D$  and  $\sigma' = \sigma$ . Then, for every  $a = (u, v) \in A(D)$ , we add an arc  $a' \in (v, u) \in A(D')$  and set  $\sigma'(a') = (\sigma(a))^{-1}$ . Finally, we remove duplicate arcs as follows. As long as there are arcs  $a, a' = (u, v) \in A(D')$  such that  $a \in A(D)$ ,  $a' \in A(D') \setminus A(D)$  and  $\sigma'(a) = \sigma'(a')$ , we remove the arc  $a'$ .*

Observe that it is possible that  $\text{Doubling}(D) = (D, \sigma)$ . In fact such graphs will be of special interest to us. The following lemma generalizes Proposition 2.1.

**Lemma 5.4.** *Let  $(D, \sigma)$  be a labeled digraph, where  $D$  is strongly connected. Then,  $D$  has a colorful walk if and only if  $\text{Doubling}(D)$  has a colorful walk.*

*Proof.* Since  $\text{Doubling}(D)$  is a supergraph of  $D$ , it follows that if  $D$  has a colorful walk then so does  $\text{Doubling}(D)$ . We now argue the converse. Let  $(D', \sigma') = \text{Doubling}(D)$  and suppose that  $D'$  has a colorful walk. Let  $v$  be an arbitrary vertex in this subgraph. We now argue that  $D$  has a colorful walk as well. More specifically, we argue that  $D$  has a  $v$ -colorful walk. By Lemma 5.3, it is sufficient to show that for every  $i \in [\ell]$ , there is a  $j \neq i$  such that  $H_D$  has a  $v_i$ - $v_j$  path. On the other hand, Lemma 5.3 also implies that for every  $\alpha \in [\ell]$ , there is a  $\beta \neq \alpha$  such that  $H_{D'}$  has a  $v_\alpha$ - $v_\beta$  path. Therefore, we fix  $\alpha \in [\ell]$  and let  $P$  be a  $v_\alpha$ - $v_\beta$  path in  $H_{D'}$  where  $\beta \neq \alpha$ . Our objective now is to demonstrate the existence of a  $\lambda \neq \alpha$  and a  $v_\alpha$ - $v_\lambda$  path in  $H_D$ .

Let  $x \in V(D')$  and  $\gamma \in [\ell]$  be such that  $x_\gamma$  is the *last* vertex in the traversal of  $P$  from  $v_\alpha$  towards  $v_\beta$  with the property that  $H_D$  also has a  $v_\alpha$ - $x_\gamma$  path  $P_1$ . If  $x = v$  and  $\gamma \neq \alpha$  then we are already done. Hence, we assume that this is not the case. More specifically,  $x_\gamma \neq v_\beta$ . However, observe that it could be the case that  $x_\gamma$  is  $v_\alpha$  itself.

Let  $y_\delta$  be the vertex that appears immediately after  $x_\gamma$  in the traversal of  $P$  starting from  $v_\alpha$ . Due to our choice of  $x_\gamma$  it must be the case that  $(y_\delta, x_\gamma) \in A(H_D)$  and there is no  $x_\gamma$ - $y_\delta$  path in  $H_D$ . However, since  $D$  is strongly connected, it follows that there is a  $\delta' \in [\ell]$  such that  $H_D$  has a  $v_\alpha$ - $y_{\delta'}$  path, where  $\delta' \neq \delta$ . Finally, from the existence of the arc  $(y_\delta, x_\gamma)$  in  $H_D$ , we infer the existence of an arc  $(y_{\delta'}, x_\rho)$ , where  $\rho \neq \gamma$ . We have thus obtained a  $v_\alpha$ - $x_\gamma$  path and a  $v_\alpha$ - $x_\rho$  path in  $H_D$ , where  $\gamma \neq \rho$ .

Since  $D$  is strongly connected, there is an  $x$ - $v$  path  $P_2$  in  $D$ . Clearly, either  $\sigma(P_2)(\gamma) \neq \alpha$  or  $\sigma(P_2)(\rho) \neq \alpha$ . We assume without loss of generality that it is the former and denote  $\sigma(P_2)(\gamma)$  by  $\lambda$ . As a result, we obtain an  $x_\gamma$ - $v_\lambda$  path in  $H_D$ . Since we already have a  $v_\alpha$ - $x_\gamma$  path in  $H_D$ , we conclude that there is a  $v_\alpha$ - $v_\lambda$  path in  $H_D$ . This completes the proof of the lemma.  $\square$

The NODE UNIQUE LABEL COVER problem (NODE ULC) is a special case of the COLORFUL WALK COVER problem where the input digraph  $D$  has the property that  $\text{Doubling}(D) = D$ . NODE ULC was introduced in the parameterized complexity setting by Chitnis et al. [10]. There are several FPT algorithms for NODE ULC parameterized by  $\ell$  and  $k$ , with varying dependencies on the two parameters [10, 28, 37, 29]. While any of these algorithms serve our purpose, we use the algorithm of Iwata et al. [29] as it matches the current best dependence on both parameters while achieving linear dependence on the input size.

**Proposition 5.1.** [29] NODE UNIQUE LABEL COVER can be solved in time  $\mathcal{O}(\ell^{2k}(m+n))$ .

It is straightforward to see that in an instance  $(D, \sigma, k)$  of NODE ULC, we can *forbid* any set of vertices to be part of the solution by simply making  $k + 1$  copies of it. That is, for a vertex  $v$ , we add new vertices  $v_1, \dots, v_k$  and for every  $i \in [k]$  and  $(v, u) \in A(D)$  we add an arc  $(v_i, u)$  and set  $\sigma((v_i, u)) = \sigma((v, u))$ . This ensures that any inclusion-wise minimal set of size at most  $k$  that covers all colorful walks in the resulting digraph must be disjoint from  $v, v_1, \dots, v_k$ . We will require this operation in the proof of Lemma 5.1 where the algorithm of Proposition 5.1 will be used as a subroutine. Hence, we reformulate this proposition as follows.

**Lemma 5.5.** *There is an algorithm that, given an instance  $(D, \sigma, k)$  of NODE UNIQUE LABEL COVER and a set  $W \subseteq V(D)$ , runs in time  $\mathcal{O}(\ell^{2k} n^{\mathcal{O}(1)})$  and either returns a solution disjoint from  $W$  or correctly concludes that no such solution exists.*

We now recall the notion of shadows from [13].

**Definition 5.6.** [13] *Let  $D$  be a digraph and  $T$  be a set of terminals. Let  $X \subseteq V(G)$  be a subset of vertices.*

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

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

The notions of shadows and ‘shadowless solutions’ have proved to be a key component of several FPT algorithms for cut-problems [13, 14, 36, 33].

**Definition 5.7.** *Let  $I = (D, \sigma, k, W)$  be an instance of RESTRICTED COLORFUL WALK COVER and let  $S$  be a solution for this instance. If the shadow of  $S$  with respect to  $W$  is empty, then we say that  $S$  is a shadowless solution.*

Combining the notion of shadowless solutions with Lemma 5.4, we make the following crucial observation that implies that if the shadow of a solution  $S$  with respect to  $W$  is empty, then the set  $S$  is a colorful walk cover even for the labeled digraph  $\text{Doubling}(D)$ .

**Observation 5.2.** *Let  $(D, \sigma, k, W)$  be an instance of RESTRICTED COLORFUL WALK COVER and  $S$  be a solution for this instance. Let  $C$  be the unique strongly connected component of  $D - S$  containing  $W$ . Then,  $\text{Doubling}(C)$  does not have a colorful walk.*

Because of Lemma 5.4, Proposition 5.1 and Observation 5.2, our objective from now on is to transform a given instance  $I$  of RESTRICTED COLORFUL WALK COVER into an instance  $I'$  such that if  $I$  is a yes-instance then  $I'$  has a shadowless solution and is a no-instance otherwise. This transformation has two steps – (a) the shadow covering step and (b) the shadow removal step. For the shadow covering step, we will use a result of Chitnis et al. [13]. Building on the work of Marx and Razgon [44] and Chitnis et al. [14], they gave a generic method to compute a set that ‘covers’ the shadow of a solution when dealing with cut-problems satisfying certain properties. In order to make this statement precise and describe how it applies in our setting, we begin with the following definition.

**Definition 5.8.** [13] *Let  $D$  be a digraph and let  $\mathcal{F} = \{F_1, \dots, F_q\}$  be a set of subgraphs of  $D$ . An  $\mathcal{F}$ -transversal is a set of vertices of  $D$  that intersects every  $F \in \mathcal{F}$ . For a set  $T \subseteq V(G)$ , we say that  $\mathcal{F}$  is  $T$ -connected if, for every  $i \in [q]$ , each vertex of  $F_i$  can reach some vertex of  $T$  by a walk completely contained in  $F_i$  and is reachable from some vertex of  $T$  by a walk completely contained in  $F_i$ .*

The notion of  $T$ -connectivity of a family  $\mathcal{F}$  is relevant to us because in an instance  $(D, \sigma, k, W)$  of RESTRICTED COLORFUL WALK COVER the set of all colorful walks is clearly  $W$ -connected. This is a consequence of the fact that  $W$  intersects all colorful walks in  $D$  and each colorful walk itself is a strongly connected subgraph of  $D$ . As a result, we can utilise the following *shadow covering* lemma of Chitnis et al.

**Proposition 5.2.** (Theorem 3.6 [13]) *Let  $D$  be a digraph and let  $T \subseteq V(D)$ . Given  $T$  and  $D$ , we can construct a set  $\{Z_1, \dots, Z_t\}$  with  $t = 2^{\mathcal{O}(k^2)} \log^2 n$  in time  $2^{\mathcal{O}(k^2)} n^{\mathcal{O}(1)}$  such that, for any set  $F$  of  $T$ -connected subgraphs, if there exists an  $\mathcal{F}$ -transversal of size at most  $k$ , then there is an  $\mathcal{F}$ -transversal  $X$  of size at most  $k$  such that for at least one  $i \in [t]$ , we have*

- $X \cap Z_i = \emptyset$ , and
- $Z_i$  covers the shadow of  $X$  with respect to  $T$ .

We use the fact that the set of colorful walks in an instance of RESTRICTED COLORFUL WALK COVER is  $W$ -connected to reformulate Proposition 5.2 as follows, so that it is easier to invoke in our context.

**Lemma 5.6.** *Let  $I = (D, \sigma, k, W)$  be an instance of RESTRICTED COLORFUL WALK COVER. Given  $I$ , we can construct a set  $\{Z_1, \dots, Z_t\}$  with  $t = 2^{\mathcal{O}(k^2)} \log^2 n$  in time  $2^{\mathcal{O}(k^2)} n^{\mathcal{O}(1)}$  such that if  $I$  is a yes-instance, then there is a colorful walk cover  $S$  of size at most  $k$  such that for at least one  $i \in [t]$ , we have*

- $S \cap Z_i = \emptyset$ , and
- $Z_i$  covers the shadow of  $S$  with respect to  $W$ .

*A set  $Z_i$  with these two properties is called a shadow cover for  $S$  with respect to  $W$ .*

## 5.2 Shadow-removal and FPT Algorithm for RESTRICTED COLORFUL WALK COVER

We now proceed to the *shadow-removal step*. This is the more problem-specific part of the approach from [13]. In this step, our objective is to *remove* the vertices in a shadow cover for a solution  $S$  in such a way that  $S$  becomes a shadowless solution for the resulting instance and we do not create new, smaller solutions for this instance. In order to achieve this, we need to define an appropriate analogue of the *torso* operation [13]. However, before we define this operation, we prove the following lemma that gives a subroutine required in the definition of this operation.

**Lemma 5.7.** *Let  $(D, \sigma)$  be a labeled digraph and let  $Z \subseteq V(D)$ . There is an algorithm that, given  $(D, \sigma)$ ,  $Z$  and a pair of (not necessarily distinct) vertices  $u, v \notin Z$ , runs in time  $\mathcal{O}(\ell^\ell) + n^{\mathcal{O}(1)}$  and computes a set  $\mathcal{X} = \{\pi_1, \dots, \pi_r\}$  of permutations of  $[\ell]$  such that for every  $\alpha, \beta \in [\ell]$ , there is a  $u$ - $v$  walk  $P$  with all internal vertices in  $Z$  such that  $\sigma(P)(\alpha) = \beta$  if and only if there is a permutation  $\pi \in \mathcal{X}$  such that  $\pi(\alpha) = \beta$ .*

*Proof.* We first construct the subgraph  $D' = D[\{u, v\} \cup Z]$  with  $\sigma'$  being the restriction of  $\sigma$  to the arcs in  $D'$ . We then construct the digraph  $H_{D', \sigma'}$ . In order to define  $\mathcal{X}$ , we will now construct another auxiliary digraph  $Q$  with vertex set  $[\ell]$  from which we will then extract the set  $\mathcal{X}$ .

For every  $i, j \in [\ell]$  we test whether there is a  $u_i$ - $v_j$  path in  $H_{D'}$  with all internal vertices in  $\mathcal{Z} = \bigcup_{v \in Z} \{v_1, \dots, v_\ell\}$ , and if such a path exists, then we add the arc  $(i, j)$ . This completes the construction of  $Q$ . Observe that  $Q$  may contain self-loops. Before we define  $\mathcal{X}$ , we prove the following property of  $Q$ . Recall that a cycle cover of a digraph is a set of vertex-disjoint cycles such that every vertex is part of some cycle.

**Claim 5.1.** *Every arc  $a \in A(Q)$  is part of a cycle cover in  $Q$ .*

*Proof.* Let  $a = (\alpha, \beta) \in A(Q)$ . By the definition of  $Q$ , it must be the case that there is a  $u_\alpha$ - $v_\beta$  path in  $H_{D', \sigma'}$  with all internal vertices in  $Z$ , which in turn implies the presence of a  $u$ - $v$  walk  $P$  in  $D$  such that all internal vertices of  $P$  lie in  $Z$  and  $\sigma(P)(\alpha) = \beta$ . Therefore the permutation  $\sigma(P)$  contains a cycle of the form  $(\dots \alpha \beta \dots)$ . As a result, the cycle cover of  $Q$  that corresponds to  $\sigma(P)$  contains the arc  $a$ . This completes the proof of the claim.  $\square$

We are now ready to define  $\mathcal{X}$ . For every arc  $a = (\alpha, \beta) \in A(Q)$ , we pick an arbitrary cycle cover in  $Q$  containing  $a$  and call it  $\mathcal{C}_a$ . By the claim above, we know that such a cycle cover exists. For every such cycle cover  $\mathcal{C}_a$ , we define the permutation  $\pi_a$  as the corresponding permutation of  $[\ell]$ . Finally, we define  $\mathcal{X} = \{\pi_a | a \in A(Q)\}$ . This completes the construction of  $\mathcal{X}$  and we now argue that it satisfies the required properties.

By Lemma 5.3, we know that for every  $i, j \in [\ell]$ , there is a  $u$ - $v$  walk  $P$  such that  $\sigma(P)(i) = j$  if and only if there is a  $u_i$ - $v_j$  path  $P'$  in  $H_{D'}$ . Furthermore, it is easy to see that  $P$  has all internal vertices in  $Z$  if and only if  $P'$  has all internal vertices in  $Z$ . Now, suppose that for some  $\alpha, \beta \in [\ell]$ , there is a  $u$ - $v$  walk  $P$  with all internal vertices in  $Z$  such that  $\sigma(P)(\alpha) = \beta$ . Then, there is a  $u_\alpha$ - $v_\beta$  path in  $H_{D'}$  with all internal vertices in  $\bigcup_{v \in Z} \{v_1, \dots, v_\ell\}$ , which by the construction of the digraph  $Q$  implies that  $(\alpha, \beta) \in A(Q)$  and hence there is a cycle in  $Q$  that contains the arc  $(\alpha, \beta)$ . Due to Claim 5.1, we conclude that there is a permutation  $\pi \in \mathcal{X}$  such that  $\pi(\alpha) = \beta$ . This completes the argument in the forward direction. The converse direction follows by retracing the above argument. Note that the required time is the time required to compute  $Q$  plus the time required to compute the cycle covers for the arcs in  $Q$ . The first part takes polynomial time and the second can be achieved by simply enumerating all cycle covers of  $Q$ . Since  $|Q| = \ell$ , the second step only requires time  $\mathcal{O}(\ell^\ell)$ , hence completing the proof of the lemma.  $\square$

We are now ready to define the labeled-torso operation.

**Definition 5.9.** *Let  $(D, \sigma)$  be a labeled digraph,  $Z \subseteq V(D)$  and for every ordered pair  $(u, v) \in (V(D) \setminus Z)^2$ , let  $\mathcal{X}_{uv}$  be the set of permutations returned by the algorithm of Lemma 5.7 on input  $(D, \sigma)$ ,  $Z$ , and the pair  $u, v$ . We let  $\mathcal{H}$  denote the set  $\{\mathcal{X}_{uv} | (u, v) \in (V(D) \setminus Z)^2\}$ . We denote by  $\text{labeled-torso}(D, Z, \mathcal{H})$  the labeled digraph  $(D', \sigma')$  obtained from  $D$  as follows.*

- Set  $D' = D, \sigma' = \sigma$ .
- Delete the set  $Z$ .
- For every ordered pair  $(u, v)$  in  $(V(D) \setminus Z)^2$ , add  $q = |\mathcal{X}_{uv}|$  arcs  $a_1^{uv}, \dots, a_q^{uv} = (u, v)$  and for each  $j \in [q]$ , set  $\sigma'(a_j^{uv}) = \pi_j$  where  $\mathcal{X}_{uv} = \{\pi_1, \dots, \pi_q\}$ .

Having defined the labeled-torso operation we proceed to show that it preserves all colorful walks.

**Lemma 5.8.** *Let  $(D, \sigma), Z, \mathcal{H}$  be as in Definition 5.9 and let  $(D', \sigma') = \text{labeled-torso}(D, Z, \mathcal{H})$ . Then, the following statements hold.*

- For every (not necessarily distinct)  $u, v \in V(D) \setminus Z$  and  $\alpha, \beta \in [\ell]$ , if there is a  $u$ - $v$  walk  $P$  in  $D$  such that  $\sigma(P)(\alpha) = \beta$  then there is a  $u$ - $v$  walk  $P'$  in  $D'$  such that  $\sigma'(P')(\alpha) = \beta$  and  $V(P') = V(P) \setminus Z$ .
- For every (not necessarily distinct)  $u, v \in V(D')$  and  $\alpha, \beta \in [\ell]$ , if there is a  $u$ - $v$  walk  $P'$  in  $D'$  such that  $\sigma'(P')(\alpha) = \beta$  then there is a  $u$ - $v$  walk  $P$  in  $D$  such that  $\sigma(P)(\alpha) = \beta$  and  $V(P) \subseteq V(P') \cup Z$ .

*Proof.* For the first statement, let  $u, v \in V(D) \setminus Z$  and  $\alpha, \beta \in [\ell]$  be such that there is a directed  $u$ - $v$  walk  $P = z_1, \dots, z_t$  in  $D$  with  $\sigma(P)(\alpha) = \beta$ . If  $P$  is also present in  $D'$ , then we are done. Suppose that this is not the case and let  $z_{i_1}$  and  $z_{i_2}$  be a pair of consecutive vertices in  $P$  that are not in  $Z$  such that  $i_1 < i_2$  and if  $i_1 + 1 < i_2$ , then the vertices  $z_{i_1+1}, \dots, z_{i_2-1}$  are all in  $Z$ . Let  $Q_1 = P[u, z_{i_1}]$ ,  $Q_2 = P[z_{i_1}, z_{i_2}]$  and  $Q_3 = P[z_{i_2}, v]$  be three subwalks of  $P$ . Let  $x = z_{i_1}$ ,  $y = z_{i_2}$  and furthermore, suppose that  $\gamma, \delta \in [\ell]$  are such that  $\sigma(Q_1)(\alpha) = \gamma$ ,  $\sigma(Q_2)(\gamma) = \delta$  and  $\sigma(Q_3)(\delta) = \beta$ .

Observe that  $Q_2$  is a walk with all internal vertices in  $Z$ . Then, by Lemma 5.7, there is a permutation  $\pi \in \mathcal{X}_{xy}$  such that  $\pi(\gamma) = \delta$  and by the definition of labeled-torso, there is an arc  $a = (x, y)$  in  $D'$  such that  $\sigma'(a)(\gamma) = \delta$ . Therefore, we replace the subwalk  $Q_2$  with this arc  $a$  and we do this for every such consecutive pair of vertices in  $P$  that are not in  $Z$  but all vertices in between them are in  $Z$ . The walk resulting from performing this replacement for every such pair is a  $u$ - $v$  walk  $P'$  in  $D'$  such that  $\sigma'(P')(\alpha) = \beta$  and  $V(P') = V(P) \setminus Z$ . This completes the argument for the first statement.

For the second statement, let  $P'$  be a directed  $u$ - $v$  walk in  $D'$  such that  $\sigma'(P')(\alpha) = \beta$ . If  $P'$  is also in  $D$ , then we are done. Suppose that this is not the case and let  $(x, y) \in A(D') \setminus A(D)$  be an arc in  $P'$ . Let  $Q_1 = P'[u, x]$ ,  $Q_2 = P'[x, y]$ ,  $Q_3 = P'[y, v]$  be three subwalks of  $P'$ , where  $Q_2$  is in fact the arc  $(x, y)$ , which by our assumption is not in  $D$ . Furthermore, as earlier, let  $\gamma, \delta \in [\ell]$  be such that  $\sigma(Q_1)(\alpha) = \gamma$ ,  $\sigma(Q_2)(\gamma) = \delta$  and  $\sigma(Q_3)(\delta) = \beta$ .

By the definition of labeled-torso it must be the case that there is a permutation  $\pi \in \mathcal{X}_{xy}$  such that  $\pi(\gamma) = \delta$ . Furthermore, by Lemma 5.7, we know that this can happen only when there is an  $x$ - $y$  walk  $Q'_2$  in  $D$  such that  $\sigma(Q'_2)(\gamma) = \delta$  and all internal vertices of  $Q'_2$  lie in  $Z$ . Therefore, we replace the arc  $(x, y)$  with the  $x$ - $y$  walk  $Q'_2$  which is contained in  $D$  and we do this for every arc in  $P'$  that is not in  $A(D)$ . The result is clearly a  $u$ - $v$  walk  $P$  in  $D$  such that  $\sigma(P)(\alpha) = \beta$  and  $V(P) \subseteq V(P') \cup Z$ . This completes the proof of the lemma.  $\square$

**Lemma 5.9.** *Let  $I = (D, \sigma, k, W)$  be an instance of RESTRICTED COLORFUL WALK COVER and let  $Z \subseteq V(D) \setminus W$  be such that if  $I$  is a yes-instance, then it has a solution  $S$  for which  $Z$  is a shadow-cover. There is an algorithm that, given  $I$  and  $Z$  runs in time  $\mathcal{O}(\ell^\ell n^{\mathcal{O}(1)})$  and returns an instance  $I' = (D', \sigma', W, k)$  such that if  $I$  is a no-instance, then  $I'$  is a no-instance and if  $I$  is a yes-instance then  $I'$  is a yes-instance with a shadowless solution.*

*Proof.* Let  $\mathcal{H}$  be the family of sets from Definition 5.9. Since  $\mathcal{H}$  can be computed by invoking the algorithm of Lemma 5.7 for every pair of vertices in  $V(D) \setminus Z$ , it follows that the time required to compute  $\mathcal{H}$  is  $\mathcal{O}(\ell^\ell n^{\mathcal{O}(1)})$ . Let  $(D', \sigma') = \text{labeled-torso}(D, Z, \mathcal{H})$  and  $I' = (D', \sigma', k, W)$ . We now argue that  $I'$  satisfies the required properties.

We first argue that if  $I$  is a no-instance then  $I'$  is a no-instance. In order to do so, we argue that any colorful walk cover of  $D'$  is also a colorful walk cover of  $D$ . Suppose that this is not the case and let  $S$  be a colorful walk cover of  $D'$  that is not a colorful walk cover of  $D$ . Then,  $D - S$  has a colorful walk  $H$ . If  $H$  is contained entirely in  $Z$ , then it contradicts our assumption that  $Z$  is a shadow-cover for  $S$ . Hence, we may assume that  $H$  has at least one vertex outside  $Z$ , call it  $h$ . Now, Lemma 5.8 implies that for every  $h$ -walk  $P$  in  $D$  and  $\alpha, \beta$  such that  $\sigma(P)(\alpha) = \beta$ , there is an  $h$ -walk  $P'$  in  $D'$  such that  $\sigma'(P')(\alpha) = \beta$  and  $V(P') = V(P) \setminus Z$ . Since  $P$  is disjoint from  $S$ , we conclude that  $P'$  is also disjoint from  $S$ . But this implies that there is an  $h$ -colorful walk in  $D'$  that is disjoint from  $S$ , a contradiction.

In the converse direction, we argue that if  $I'$  is a no-instance then  $I$  is a no-instance. In order to do so, we argue that any colorful walk cover of  $D$  disjoint from  $Z$  is a colorful walk cover of  $D'$  disjoint from  $S$ . Suppose that this is not the case and let  $S$  be a colorful walk cover of  $D$  disjoint from  $Z$  that is not a colorful walk cover of  $D'$ . Then, there is a vertex  $h \in V(D')$  and a subgraph  $H$  that is an  $h$ -colorful walk in  $D' - S$ .

By Lemma 5.8, we know that for every  $u, v, \alpha, \beta$  and a  $u$ - $v$  walk  $P'$  in  $D'$  such that  $\sigma'(P')(\alpha) = \beta$ , there is a  $u$ - $v$  walk  $P$  in  $D$  such that  $\sigma(P)(\alpha) = \beta$  and  $V(P) \subseteq V(P') \cup Z$ . Since  $P'$  is disjoint from  $S$  and  $S$  is disjoint from  $Z$ , it follows that  $P$  is also disjoint from  $S$ . As a result, we infer the presence of an  $h$ -colorful walk in  $D - S$  as well, a contradiction.

Finally, observe that due to Lemma 5.8, every pair of vertices in the same strongly connected component as  $W$  in  $D - S$  remain in the same strongly connected component as  $W$  in  $D' - S$ . Furthermore,  $Z$  covers all vertices in the shadow of  $S$  with respect to  $W$  in  $D$  and  $V(D') = V(D) \setminus Z$ . Therefore, the shadow of  $S$  with respect to  $W$  in  $D'$  is empty. This completes the proof of the lemma.  $\square$

We are now ready to complete the proof of Lemma 5.1.

**Lemma 5.1.** RESTRICTED CWC has an algorithm running in time  $\ell^{\mathcal{O}(k+\ell)} 2^{\mathcal{O}(k^2)} n^{\mathcal{O}(1)}$ .

*Proof.* Let  $I = (D, \sigma, k, W)$  be the given instance of RESTRICTED COLORFUL WALK COVER. We first execute the algorithm of Lemma 5.6 on this instance and obtain sets  $\{Z_1, \dots, Z_t\}$ , where  $t = 2^{\mathcal{O}(k^2) \log^2 n}$ . For every  $i \in [t]$ , we execute the algorithm of Lemma 5.9 on input  $I$  and  $Z_i$  to obtain the instance  $I'_i = (D'_i, \sigma'_i, k, W)$ . Finally, for each  $i \in [t]$ , we execute the NODE UNIQUE LABEL COVER algorithm of Lemma 5.5 on input  $I''_i = (\text{Doubling}(D'_i, \sigma'_i), k)$  and the set  $W$  to either compute a colorful walk cover of size at most  $k$  disjoint from  $W$  or correctly conclude that no such set exists. Finally, if for any  $i \in [t]$ , the solution to  $I''_i$  is not NO, then we return the computed set as the solution for the given instance of RESTRICTED COLORFUL WALK COVER.

The correctness and the claimed bound on the running time of this algorithm follow from those of Lemma 5.6, Lemma 5.9, Lemma 5.5 and Observation 5.2. This completes the proof of the lemma.  $\square$

### 5.3 The Parameterized Approximation for COLORFUL WALK COVER

We are finally ready to complete the proof of Theorem 5. For the sake of completeness, we restate it here.

**Theorem 5.** COLORFUL WALK COVER admits an  $\ell^{\mathcal{O}(k+\ell)} 2^{\mathcal{O}(k^2)} n^{\mathcal{O}(1)}$  time FPT-approximation algorithm with approximation ratio 2.

At the highest level, our algorithm utilises the *iterative compression* technique introduced by Reed, Smith and Vetta [48] in order to prove the fixed-parameter tractability of the ODD CYCLE TRANSVERSAL problem on undirected graphs. It has subsequently gone on to become a fundamental tool in the FPT-algorithmist's toolbox. We now proceed to provide a description of the highest level of our algorithm that uses this technique.

Given an instance  $(D, \sigma, k)$  of COLORFUL WALK COVER, where  $V(D) = \{v_1, \dots, v_n\}$ , we define a labeled graph  $(D_i, \sigma_i)$  where  $V_i = \{v_1, \dots, v_i\}$ ,  $D_i = D[V_i]$  and  $\sigma_i$  is the restriction of  $\sigma$  to  $V_i$ . We iterate through the instances  $(D_i, \sigma_i, k)$  starting from  $i = 2k + 1$  and for the  $i^{\text{th}}$  instance, with the help of a *known* solution  $\hat{S}_i$  of size at most  $2k + 1$  we either correctly conclude that the  $i^{\text{th}}$  instance has no colorful walk cover of size at most  $k$  or try to find a colorful walk cover  $S_i$  of size at most  $2k$ , i.e, a 2-approximation. This problem, which is known as the *compression* problem is formally defined as follows.

COLORFUL WALK COVER COMPRESSION

**Parameter:**  $k, \ell$

**Input:**  $(D, \sigma, k, \hat{S})$  where  $\hat{S}$ , a colorful walk cover of size at most  $2k + 1$ .

**Question:** Does there exist a colorful walk cover of size at most  $k$  for this instance?

Our algorithm for the COLORFUL WALK COVER problem comprises of ‘solving’ at most  $n$  instances of the COLORFUL WALK COVER COMPRESSION problem. Henceforth, in this context, we use ‘solving’ to also mean obtaining a 2-approximate solution. Let  $I_i = (D_i, \sigma, k, \hat{S}_i)$  be the  $i^{\text{th}}$  instance of COLORFUL WALK COVER COMPRESSION. Clearly, the set  $V_{2k+1}$  is a solution of size at most  $2k+1$  for the instance  $I_{2k+1}$ . It is also easy to see that if  $S_{i-1}$  is a colorful walk cover of size at most  $2k$  for instance  $I_{i-1}$ , then the set  $S_{i-1} \cup \{v_i\}$  is a colorful walk cover of size at most  $2k+1$  for the instance  $I_i$ . We use these two observations to initiate the iteration with the instance  $(D_{2k+1}, \sigma, k, \hat{S}_{2k+1} = V_{2k+1})$  and either compute a colorful walk cover of size at most  $2k$  for this instance or correctly conclude that there is no colorful walk cover of size at most  $k$ . If there is such a solution  $S_{2k+1}$ , then we set  $\hat{S}_{2k+2} = S_{2k+1} \cup \{v_{2k+2}\}$  and try to compute a colorful walk cover of size at most  $2k$  for the instance  $I_{k+2}$  and so on. If, on the other hand during any iteration, the corresponding instance is found to *not* have a colorful walk cover of size at most  $k$ , then it implies that the original instance is a NO instance. Since the only way we proceed in the iteration is by computing a 2-approximate colorful walk cover  $S_i$  for the instance  $I_i$  of COLORFUL WALK COVER COMPRESSION, the required 2-approximate colorful walk cover for the original input instance will be  $S_n$ . Since there can be at most  $n$  iterations, the total time taken is bounded by  $n$  times the time required to solve the COLORFUL WALK COVER COMPRESSION problem. We now discuss how to solve the COLORFUL WALK COVER COMPRESSION problem by reducing it to a bounded number of instances of the RESTRICTED COLORFUL WALK COVER problem. However, before we proceed, we need the following definition and proposition (see [16]).

**Definition 5.10.** Let  $D$  be a digraph and  $\mathcal{X} = \{X_1, \dots, X_r\}$  be a set of disjoint vertex sets of  $D$ . A set  $S \subseteq V(D) \setminus \bigcup_{i \in [r]} X_i$  is called an  $\mathcal{X}$ -skew separator if for every  $1 \leq i < j \leq r$ , there is no directed  $X_i$ - $X_j$  path in  $D - S$ .

**Proposition 5.3.** [16] There is an algorithm that, given a digraph  $D$ , a set  $\mathcal{X} = \{X_1, \dots, X_r\}$  of disjoint vertex sets and an integer  $k$ , runs in time  $\mathcal{O}(4^k n^{\mathcal{O}(1)})$  and either returns an  $\mathcal{X}$ -skew separator of size at most  $k$  or correctly concludes that one does not exist.

**Lemma 5.10.** Let  $(D, \sigma, k, \hat{S})$  be an instance of COLORFUL WALK COVER COMPRESSION and let  $S$  be a solution for this instance. There exists an ordered partition  $\mathcal{W} = \{W_1, \dots, W_r\}$  of  $\hat{S} \setminus S$  such that  $S$  is a  $\mathcal{W}$ -skew separator in  $D$ .

*Proof.* Let  $\mathcal{C} = \{C_1, \dots, C_t\}$  be the set of strongly connected components of  $D - S$  such that for every  $1 \leq i < j \leq t$ , there is no  $C_i$ - $C_j$  path in  $D$ . Let  $\mathcal{C}' = \{C_{i_1}, \dots, C_{i_r}\}$  be the subset of  $\mathcal{C}$  comprising strongly connected components intersecting  $\hat{S} \setminus S$ , where  $i_j < i_{j'}$  for every  $1 \leq j < j' \leq r$ . For each  $j \in [r]$ , let  $W_j = (\hat{S} \setminus S) \cap C_j$ . From the definitions of  $\mathcal{C}$  and  $\mathcal{C}'$ , it follows that there is no  $W_i$ - $W_j$  path in  $D - S$  for any  $1 \leq i < j \leq r$ , implying that  $S$  is a  $\mathcal{W}$ -skew separator, where  $\mathcal{W} = \{W_1, \dots, W_r\}$ . This completes the proof of the lemma.  $\square$

We refer to the unique partition  $\mathcal{W}$  in the proof of the above lemma, as the partition of  $\hat{S} \setminus S$  that *respects*  $S$ .

**Definition 5.11.** Let  $I = (D, \sigma, k, \hat{S})$  be an instance of COLORFUL WALK COVER COMPRESSION and let  $\mathcal{W} = \{W_1, \dots, W_r\}$  be a partition of  $\hat{S}$  such that  $D$  has a  $\mathcal{W}$ -skew separator of size 0. Suppose that if  $I$  is a yes-instance then  $\mathcal{W}$  is the partition respecting a solution for  $I$  disjoint from  $\hat{S}$ . Then  $I$  is called a  $\mathcal{W}$ -nice instance.

**Lemma 5.11.** There is an algorithm that, given an instance  $I = (D, \sigma, k, \hat{S})$  of COLORFUL WALK COVER COMPRESSION and a partition  $\mathcal{W} = \{W_1, \dots, W_r\}$  of  $\hat{S}$  such that  $I$  is a  $\mathcal{W}$ -nice instance, runs in time  $\ell^{\mathcal{O}(k+\ell)} 2^{\mathcal{O}(k^2)} n^{\mathcal{O}(1)}$  and either returns a colorful walk cover of size at most  $k$  or correctly concludes that one does not exist.

*Proof.* Let  $S$  be a solution disjoint from  $\hat{S}$  such that  $\mathcal{W}$  is the partition that respects  $S$ . For each  $i \in [r]$ , let  $C_i$  denote the unique strongly connected component of  $D$  that contains the set  $W_i$ . Since  $I$  is a  $\mathcal{W}$ -nice instance, it has a  $\mathcal{W}$ -skew separator of size 0. Therefore, for  $1 \leq i < j \leq r$ ,  $C_i$  and  $C_j$  are distinct. For every  $i \in [r]$ , we now define the labeled digraph  $(D_i, \sigma_i) = (D[C_i], \sigma|_{C_i})$  and the instance  $I_i = (D_i, \sigma_i, k, W_i)$  of RESTRICTED COLORFUL WALK COVER.

We execute the algorithm of Lemma 5.1 for each  $i \in [r]$  and let  $L_i$  denote the result of the execution on the instance  $I_i$ , where  $L_i$  can either denote NO or a smallest solution for the instance  $I_i$ . If for any  $i \in [r]$ ,  $L_i$  is NO, then we return that  $I$  is a no-instance. This is correct because  $I_i$  is a sub-instance of  $I$ . On the other hand, suppose that for each  $i \in [r]$ ,  $L_i$  denotes a vertex set that we know is a *smallest* colorful walk cover of  $D_i$  of the required kind. Since the digraphs  $D_1, \dots, D_r$  are vertex-disjoint and every colorful walk is contained in one of these digraphs, we conclude that  $S' = \bigcup_{i \in [r]} S_i$  is a smallest colorful walk cover for the instance  $I$ . Therefore, if  $S'$  is larger than  $k$  then we return NO and otherwise we return  $S'$ . The running time of this algorithm is dominated by the time required for at most  $2k + 1$  invocations of the algorithm of Lemma 5.1, proving the stated bound on the running time. This completes the proof of the lemma.  $\square$

**Definition 5.12.** Let  $(D, \sigma)$  be a labeled digraph. A consistent labeling of  $D$  is a function  $\Gamma : V(D) \rightarrow [\ell]$  such that for every arc  $a = (u, v) \in A(D)$ ,  $\sigma(a)(\Gamma(u)) = \Gamma(v)$ . For a set  $X \subseteq V(D)$  and function  $\chi : X \rightarrow [\ell]$ , we say that  $\chi$  is an extendible consistent labeling of  $D$  if  $D$  has a consistent labeling  $\Gamma$  such that  $\Gamma|_X = \chi$ .

**Proposition 5.4.** [10] Let  $(D, \sigma)$  be a strongly connected labeled digraph such that  $\text{Doubling}(D) = D$ . Then,  $D$  has a consistent labeling if and only if it has no colorful walks.

Combining Proposition 5.4 and Lemma 5.4, we make the following observation.

**Observation 5.3.** A strongly connected labeled digraph has a consistent labeling if and only if it has no colorful walks.

The above observation implies that every strongly connected component of  $D - S$  has a consistent labeling. We now define the operation of *bundling* a set of vertices as follows.

**Definition 5.13.** Let  $(D, \sigma)$  be a labeled digraph. Let  $X \subseteq V(D)$  and  $\Gamma : X \rightarrow [\ell]$ . We denote by  $\text{Bundle}(D, \sigma, X, \Gamma)$  the labeled digraph  $(D', \sigma')$  obtained from  $D$  as follows. Initially,  $D' = D$ ,  $\sigma' = \sigma$ . For every pair  $x_1, x_2 \in X$ , we pick an arbitrary permutation  $\pi \in S_\ell$  such that  $\pi(\Gamma(x_1)) = \Gamma(x_2)$  and we add an arc  $a = (x_1, x_2)$  with  $\sigma'(a) = \pi$ .

Note that this operation is essentially the same as identifying the vertices of  $X$  to get a single new vertex and then updating the labels on the arcs adjacent to the resulting new vertex in a certain way specified by the function  $\Gamma$ . However, we define it in this way because it simplifies the presentation in the rest of the section.

**Lemma 5.12.** Let  $(D, \sigma)$  be a strongly connected labeled digraph with no colorful walks and let  $\Gamma : V(D) \rightarrow [\ell]$  be a consistent labeling. Then, for any  $X \subseteq V(D)$ ,

- $\Gamma$  is a consistent labeling for the labeled digraph  $D' = \text{Bundle}(D, \sigma, X, \Gamma|_X)$  and
- $D'$  does not contain a colorful walk.

*Proof.* The first statement is a simple consequence of the fact that  $\Gamma$  is already a consistent labeling of  $D$  and the newly added arcs clearly do not violate the condition required for  $\Gamma$  to be consistent. The second statement of the lemma follows from the first statement and Observation 5.3.  $\square$

We are now ready to present our algorithm that solves the COLORFUL WALK COVER COMPRESSION problem. That is, an algorithm that, if the given instance is not a NO instance, returns a colorful walk cover whose size is at most twice the given budget.

**Lemma 5.13.** *There is an algorithm that, given an instance  $(D, \sigma, k, \hat{S})$  of COLORFUL WALK COVER COMPRESSION, runs in time  $\ell^{\mathcal{O}(k)} 2^{\mathcal{O}(k^2)} n^{\mathcal{O}(1)}$  and either computes a colorful walk cover of size at most  $2k$  or correctly concludes that there is no colorful walk cover of size at most  $k$ .*

*Proof.* Let  $S$  be a solution for the given instance of COLORFUL WALK COVER COMPRESSION and let  $Y = S \cap \hat{S}$ . Let  $\mathcal{W}_Y = \{W_1, \dots, W_r\}$  be an ordered partition of  $\hat{S} \setminus Y$  that respects  $S$ .

We first guess the set  $Y$  and the partition  $\mathcal{W}_Y$ . Furthermore, we guess a function  $\Gamma : \hat{S} \setminus S \rightarrow [\ell]$  such that for every  $i \in [r]$ , the restriction  $\Gamma|_{W_i}$  is an extendible consistent labeling of the strongly connected component of  $D - S$  containing  $W_i$ . Due to Observation 5.3, such a  $\Gamma$  must exist. Furthermore, since  $|\hat{S}|$  is bounded by  $2k + 1$ , there are  $\ell^{\mathcal{O}(k)}$  choices for  $\Gamma$ . We now construct a new digraph by ‘bundling’ each set in  $\mathcal{W}$ . This is done as follows. For each  $i \in [r]$ , we define the graph  $(D_i, \sigma_i) = \text{Bundle}(D_{i-1}, \sigma_{i-1}, \Gamma_{W_i}, W_i)$ , where  $(D_0, \sigma_0) = (D, \sigma)$ . Clearly, the strongly connected components of  $D_r - S$  are the same as those of  $D - S$  and by Lemma 5.12, it follows that  $\Gamma|_{W_i}$  is still an extendible consistent labeling for the strongly connected component of  $D_r - S$  containing  $W_i$ . Furthermore, for *any* set  $X$  disjoint from  $\hat{S} \setminus S$ , for every  $i \in [r]$ , the vertices in  $W_i$  remain in the same strongly connected component of  $D_r - X$ .

We now execute the algorithm of Lemma 5.11 to compute a  $\mathcal{W}$ -skew separator of size at most  $k$ . If no such separator exists, then by Lemma 5.10, we may correctly conclude that the instance  $I$  is a no-instance and hence we return the same. On the other hand, let  $X$  be a  $\mathcal{W}$ -skew separator of size at most  $k$  and let  $D' = D_r - X$  with  $\sigma'$  being the associated labeling function. Observe that there is a  $\mathcal{W}$ -skew separator of size 0 in  $D'$  and  $D'$  now has a colorful walk cover  $S \setminus X$  of size at most  $k$  such that the partition  $\mathcal{W}$  respects  $S \setminus X$ .

We now construct the instance  $(D', \sigma', k, \hat{S} \setminus (X \cup Y))$  of COLORFUL WALK COVER COMPRESSION that as we have already argued, is a  $\mathcal{W}$ -nice instance. We then execute the algorithm of Lemma 5.11 to compute a colorful walk cover  $Z$  of size at most  $k$  for  $D'$ . If no such set exists, then  $I$  is a no-instance and we return the same. Otherwise, the set  $X \cup Z$  is a colorful walk cover for  $D$  of size at most  $2k$ . Hence, we return  $X \cup Z$ . This completes the description of the algorithm. The bound on the running time follows from that of Lemma 5.11 and the fact the number of invocations of the algorithm of this lemma is bounded by the product of the number of choices for  $Y$ ,  $\Gamma$  and  $\mathcal{W}$ . Since this is bounded by  $\ell^{\mathcal{O}(k)} 2^{\mathcal{O}(k \log k)}$ , the running time follows and this completes the proof of the lemma.  $\square$

## 6 Conclusion

Our results on DIRECTED ODD CYCLE TRANSVERSAL raise a few natural questions.

- The first question is whether one can improve on the approximation factor of 2 in Theorem 2 or strengthen the inapproximability result in Theorem 3 to show that even such an improvement is unlikely.
- Secondly, although Theorem 1 implies that DOCT is unlikely to have a kernel of *any size*, our FPT-approximation algorithm implies that DOCT does have a 2-approximate kernel of exponential size (see Proposition 3.2, [35]). Therefore, an exciting new challenge related to DOCT is to determine whether it has a  $c$ -approximate kernel of *polynomial size* for some constant  $c$  and if so, to find the smallest such constant. Note that Theorem 3 also rules out a  $(1 + \epsilon)$ -approximate kernel (for some  $\epsilon > 0$ ) of *any size* for DOCT.

We conclude by pointing out that the parameterized complexity of the DIRECTED MULTICUT problem where the number of terminal pairs is 3, remains open. As was the case for DOCT, it is quite likely that an FPT algorithm or a W-hardness proof for this problem would require new insights into the structure of directed cuts.

**Acknowledgements.** The authors would like to thank Michał Włodarczyk for enlightening discussions on the DOCT problem.

## References

- [1] A. AGARWAL, M. CHARIKAR, K. MAKARYCHEV, AND Y. MAKARYCHEV,  *$O(\sqrt{\log n})$  approximation algorithms for min uncut, min 2cnf deletion, and directed cut problems*, in Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005, 2005, pp. 573–581. [1](#)
- [2] N. ALON, O. SCHWARTZ, AND A. SHAPIRA, *An elementary construction of constant-degree expanders*, *Combinatorics, Probability & Computing*, 17 (2008), pp. 319–327. [35](#)
- [3] J. BANG-JENSEN AND G. GUTIN, *Digraphs - theory, algorithms and applications*, Springer, 2002. [6](#)
- [4] A. BHATTACHARYYA, S. GHOSHAL, S. KARTHIK C., AND P. MANURANGSI, *Parameterized Intractability of Even Set and Shortest Vector Problem from Gap-ETH*, ArXiv e-prints, (2018). [2](#)
- [5] A. BHATTACHARYYA, S. GHOSHAL, KARTHIK C. S., AND P. MANURANGSI, *Parameterized intractability of even set and shortest vector problem from gap-eth*, in 45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic, 2018, pp. 17:1–17:15. [2](#)
- [6] P. CHALERMSEOK, M. CYGAN, G. KORTSARZ, B. LAEKHANUKIT, P. MANURANGSI, D. NANONGKAI, AND L. TREVISAN, *From gap-eth to fpt-inapproximability: Clique, dominating set, and more*, in 58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017, 2017, pp. 743–754. [2](#), [32](#)
- [7] J. CHEN, Y. LIU, AND S. LU, *An improved parameterized algorithm for the minimum node multiway cut problem*, *Algorithmica*, 55 (2009), pp. 1–13. [3](#)
- [8] J. CHEN, Y. LIU, S. LU, B. O’SULLIVAN, AND I. RAZGON, *A fixed-parameter algorithm for the directed feedback vertex set problem*, *J. ACM*, 55 (2008). [1](#), [3](#), [4](#)
- [9] R. CHITNIS, *Directed Graphs: Fixed-Parameter Tractability & Beyond*, PhD thesis, University of Maryland, 2014. [1](#)
- [10] R. CHITNIS, M. CYGAN, M. HAJIAGHAYI, M. PILIPCZUK, AND M. PILIPCZUK, *Designing FPT algorithms for cut problems using randomized contractions*, *SIAM J. Comput.*, 45 (2016), pp. 1171–1229. [1](#), [38](#), [41](#), [48](#)
- [11] R. CHITNIS, A. E. FELDMANN, AND P. MANURANGSI, *Parameterized approximation algorithms for bidirected steiner network problems*, in 26th Annual European Symposium on Algorithms, ESA 2018, August 20-22, 2018, Helsinki, Finland, 2018, pp. 20:1–20:16. [2](#)
- [12] R. CHITNIS AND M. T. HAJIAGHAYI, *Shadowless solutions for fixed-parameter tractability of directed graphs*, in *Encyclopedia of Algorithms*, 2016, pp. 1963–1966. [1](#), [5](#)

- [13] R. H. CHITNIS, M. CYGAN, M. T. HAJIAGHAYI, AND D. MARX, *Directed subset feedback vertex set is fixed-parameter tractable*, ACM Transactions on Algorithms, 11 (2015), p. 28. [1](#), [3](#), [5](#), [42](#), [43](#)
- [14] R. H. CHITNIS, M. HAJIAGHAYI, AND D. MARX, *Fixed-parameter tractability of directed multiway cut parameterized by the size of the cutset*, SIAM J. Comput., 42 (2013), pp. 1674–1696. [3](#), [5](#), [42](#)
- [15] H. CHOI, K. NAKAJIMA, AND C. S. RIM, *Graph bipartization and via minimization*, SIAM J. Discrete Math., 2 (1989), pp. 38–47. [1](#)
- [16] M. CYGAN, F. V. FOMIN, L. KOWALIK, D. LOKSHTANOV, D. MARX, M. PILIPCZUK, M. PILIPCZUK, AND S. SAURABH, *Parameterized Algorithms*, Springer, 2015. [1](#), [3](#), [4](#), [6](#), [47](#)
- [17] M. CYGAN, M. PILIPCZUK, M. PILIPCZUK, AND J. O. WOJTASZCZYK, *Subset feedback vertex set is fixed-parameter tractable*, SIAM J. Discrete Math., 27 (2013), pp. 290–309. [3](#)
- [18] E. D. DEMAINE, G. GUTIN, D. MARX, AND U. STEGE, *07281 open problems – structure theory and FPT algorithmics for graphs, digraphs and hypergraphs*, in Structure Theory and FPT Algorithmics for Graphs, Digraphs and Hypergraphs, 08.07. - 13.07.2007, 2007. [1](#)
- [19] R. DIESTEL, *Graph Theory*, Springer-Verlag, Heidelberg, 4th ed., 2010. [6](#)
- [20] I. DINUR, *Mildly exponential reduction from gap 3sat to polynomial-gap label-cover*, Electronic Colloquium on Computational Complexity (ECCC), 23 (2016), p. 128. [2](#), [33](#)
- [21] R. G. DOWNEY AND M. R. FELLOWS, *Fundamentals of Parameterized Complexity*, Texts in Computer Science, Springer, 2013. [1](#), [4](#), [6](#)
- [22] G. EVEN, J. NAOR, B. SCHIEBER, AND M. SUDAN, *Approximating minimum feedback sets and multicuts in directed graphs*, Algorithmica, 20 (1998), pp. 151–174. [1](#)
- [23] J. FLUM AND M. GROHE, *Parameterized Complexity Theory*, Texts in Theoretical Computer Science. An EATCS Series, Springer-Verlag, Berlin, 2006. [6](#)
- [24] V. GURUSWAMI AND E. LEE, *Simple proof of hardness of feedback vertex set*, Theory of Computing, 12 (2016), pp. 1–11. [1](#)
- [25] R. IMPAGLIAZZO AND R. PATURI, *On the complexity of  $k$ -sat*, J. Comput. Syst. Sci., 62 (2001), pp. 367–375. [2](#)
- [26] R. IMPAGLIAZZO, R. PATURI, AND F. ZANE, *Which problems have strongly exponential complexity?*, J. Comput. Syst. Sci., 63 (2001), pp. 512–530. [2](#)
- [27] Y. IWATA, K. OKA, AND Y. YOSHIDA, *Linear-time FPT algorithms via network flow*, in SODA, 2014, pp. 1749–1761. [5](#)
- [28] Y. IWATA, M. WAHLSTRÖM, AND Y. YOSHIDA, *Half-integrality, lp-branching, and FPT algorithms*, SIAM J. Comput., 45 (2016), pp. 1377–1411. [1](#), [41](#)
- [29] Y. IWATA, Y. YAMAGUCHI, AND Y. YOSHIDA, *0/1/all csps, half-integral  $a$ -path packing, and linear-time FPT algorithms*, in 59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018, 2018, pp. 462–473. [41](#)
- [30] R. M. KARP, *Reducibility among combinatorial problems*, in Proceedings of a symposium on the Complexity of Computer Computations, held March 20-22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York., 1972, pp. 85–103. [1](#)

- [31] S. KHOT AND N. K. VISHNOI, *The unique games conjecture, integrality gap for cut problems and embeddability of negative-type metrics into  $\ell_1$* , J. ACM, 62 (2015), pp. 8:1–8:39. [1](#), [39](#)
- [32] S. KRATSCH, M. PILIPCZUK, M. PILIPCZUK, AND M. WAHLSTRÖM, *Fixed-parameter tractability of multicut in directed acyclic graphs*, SIAM J. Discrete Math., 29 (2015), pp. 122–144. [3](#)
- [33] D. LOKSHTANOV AND D. MARX, *Clustering with local restrictions*, Inf. Comput., 222 (2013), pp. 278–292. [3](#), [42](#)
- [34] D. LOKSHTANOV, N. S. NARAYANASWAMY, V. RAMAN, M. S. RAMANUJAN, AND S. SAURABH, *Faster parameterized algorithms using linear programming*, ACM Trans. Algorithms, 11 (2014), pp. 15:1–15:31. [5](#)
- [35] D. LOKSHTANOV, F. PANOLAN, M. S. RAMANUJAN, AND S. SAURABH, *Lossy kernelization*, in Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017, 2017, pp. 224–237. [49](#)
- [36] D. LOKSHTANOV AND M. S. RAMANUJAN, *Parameterized tractability of multiway cut with parity constraints*, in Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part I, 2012, pp. 750–761. [3](#), [42](#)
- [37] D. LOKSHTANOV, M. S. RAMANUJAN, AND S. SAURABH, *A linear-time parameterized algorithm for node unique label cover*, in 25th Annual European Symposium on Algorithms, ESA 2017, September 4-6, 2017, Vienna, Austria, 2017, pp. 57:1–57:15. [1](#), [3](#), [40](#), [41](#)
- [38] ———, *Linear time parameterized algorithms for subset feedback vertex set*, ACM Trans. Algorithms, 14 (2018), pp. 7:1–7:37. [3](#)
- [39] P. MANURANGSI AND P. RAGHAVENDRA, *A birthday repetition theorem and complexity of approximating dense csps*, in 44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland, vol. 80 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017, pp. 78:1–78:15. [2](#), [33](#)
- [40] D. MARX, *Parameterized graph separation problems*, Theoret. Comput. Sci., 351 (2006), pp. 394–406. [3](#)
- [41] D. MARX, *Can you beat treewidth?*, Theory of Computing, 6 (2010), pp. 85–112. [2](#), [7](#)
- [42] D. MARX, *What’s next? future directions in parameterized complexity*, in The Multivariate Algorithmic Revolution and Beyond - Essays Dedicated to Michael R. Fellows on the Occasion of His 60th Birthday, vol. 7370 of Lecture Notes in Computer Science, 2012, pp. 469–496. [1](#)
- [43] ———, *Some open problems in parameterized complexity (slides, dagstuhl seminar 17041)* <http://www.cs.bme.hu/~dmarx/papers/marx-dagstuhl2017-open.pdf>, 2017. [1](#)
- [44] D. MARX AND I. RAZGON, *Fixed-parameter tractability of multicut parameterized by the size of the cutset*, SIAM J. Comput., 43 (2014), pp. 355–388. [3](#), [5](#), [42](#)
- [45] R. NIEDERMEIER, *Invitation to Fixed-Parameter Algorithms*, vol. 31 of Oxford Lecture Series in Mathematics and its Applications, Oxford University Press, Oxford, 2006. [6](#)

- [46] M. PILIPCZUK AND M. WAHLSTRÖM, *Directed multicut is  $W[1]$ -hard, even for four terminal pairs*, TOCT, 10 (2018), pp. 13:1–13:18. [4](#), [7](#)
- [47] M. S. RAMANUJAN AND S. SAURABH, *Linear-time parameterized algorithms via skew-symmetric multicut*, ACM Trans. Algorithms, 13 (2017), pp. 46:1–46:25. [5](#)
- [48] B. A. REED, K. SMITH, AND A. VETTA, *Finding odd cycle transversals*, Oper. Res. Lett., 32 (2004), pp. 299–301. [1](#), [5](#), [46](#)