COUNTING HOMOMORPHISMS TO K_4-MINOR-FREE GRAPHS, MODULO 2*

JACOB FOCKE[†], LESLIE ANN GOLDBERG[†], MARC ROTH[‡], AND STANISLAV ŽIVNÝ[†]

Abstract. We study the problem of computing the parity of the number of homomorphisms from an input graph G to a fixed graph H. Faben and Jerrum [ToC'15] introduced an explicit criterion on the graph H and conjectured that, if satisfied, the problem is solvable in polynomial time and, otherwise, the problem is complete for the complexity class $\oplus$P of parity problems. We verify their conjecture for all graphs H that exclude the complete graph on 4 vertices as a minor. Further, we rule out the existence of a subexponential-time algorithm for the $\oplus$P-complete cases, assuming the randomised Exponential Time Hypothesis. Our proofs introduce a novel method of deriving hardness from globally defined substructures of the fixed graph H. Using this, we subsume all prior progress towards resolving the conjecture (Faben and Jerrum [ToC'15]; Göbel, Goldberg and Richerby [ToCT'14,'16]). As special cases, our machinery also yields a proof of the conjecture for graphs with maximum degree at most 3, as well as a full classification for the problem of counting list homomorphisms, modulo 2.

Key words. Counting modulo 2, Counting complexity, Graph homomorphisms, Parity complexity dichotomy

AMS subject classifications. 68Q17, 68Q25

1. Introduction. A homomorphism from a graph G to a graph H is a map h from $V(G)$ to $V(H)$ that preserves edges in the sense that, for every edge $\{u, v\}$ of G, the image $\{h(u), h(v)\}$ is an edge of H. Many combinatorial structures can be modelled using graph homomorphisms. For this reason, graph homomorphisms are ubiquitous in both classical and modern-day complexity theory with applications in areas such as constraint satisfaction problems [24], evaluations of spin systems in statistical physics [2, 1], database theory [31, 25], and parameterised algorithms [4, 34]. The computational problems of finding and counting homomorphisms are therefore amongst the most well-studied computational problems; the analysis of their complexity dates back to the intractability result for computing the chromatic number, one of Karp's original 21 NP-complete problems [28]. More recent work builds on Hell and Nešetřil's celebrated dichotomy theorem [26], which shows that determining whether an input graph G has a homomorphism to a fixed graph H is polynomial-time solvable if H is bipartite, or if H has a self-loop. For any other graph H, they show that the problem is NP-complete.

This paper focusses on the problem of counting homomorphisms. Applications of this problem are discussed in [1]. The complexity of the problem has been the focus of much research (see, for example, [9, 16, 17, 30, 3]).[1]

*A preliminary version of this paper (without the proofs) appeared in the proceedings of SODA 2021. The research leading to these results has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 714532). Jacob Focke has received funding from the Engineering and Physical Sciences Research Council (grant ref: EP/M508111/1). Stanislav Živný was supported by a Royal Society University Research Fellowship. The paper reflects only the authors' views and not the views of the ERC or the European Commission. The European Union is not liable for any use that may be made of the information contained therein.

[†]Department of Computer Science, University of Oxford

[‡]Merton College, University of Oxford

[1]There is also a huge literature on generalisations of this problem such as counting *weighted* homomorphisms (computing partition functions of spin systems or holant problems), counting homomorphisms to *directed* graphs, counting partition functions of constraint satisfaction problems,

The complexity of counting homomorphisms was initiated by Dyer and Greenhill [9], who gave a complete dichotomy theorem. The complexity of counting the homomorphisms from an input graph G to a fixed graph H is polynomial-time solvable if every component of H is either a complete bipartite graph with no self-loops or a complete graph in which every vertex has a self-loop. For any other graph H, they show that the problem is #P-complete.

Given that (exactly) counting the homomorphisms to H is #P-complete for almost every graph H, research has focussed on restrictions of the problem. Instead of determining the exact number of homomorphisms from G to H, compute an approximation to this number [16, 17, 30], or determine whether it is odd or even [11, 12, 19, 18], or determine its value modulo any prime p [20, 29]. Alternatively, consider the parameterised complexity [14]. For example, the problem can be studied when the input G is assumed to have bounded treewidth [6] or when H has a bounded treewidth, for example when H is a tree [12, 20, 29, 21].

Restricting the input G to have bounded treewidth makes counting homomorphisms tractable — given this restriction, the problem is solvable in polynomial time for any fixed H [6, Corollary 5.1]. Restricting the fixed target graph H to have bounded treewidth leads to a more nuanced complexity classification, even for treewidth 1 (when H is a tree). For example, the complexity of approximately counting homomorphisms to a tree H has still not been fully resolved, and it is known that different trees lead to vastly different complexities. For example, approximately counting homomorphisms to the very simple tree that is a path of length 3 is equivalent to #BIS, which is the canonical open problem in approximate counting [10]. Moreover, [21] shows that for every integer $q \geq 3$ there is a tree J_q such that approximately counting homomorphisms to J_q is equivalent to classic problem of approximating the partition function of the q-state Potts model from statistical physics. Also, it shows that there are trees H such that approximately counting homomorphisms to H is NP-hard.

1.1. Counting modulo 2 and Past Work. Faben and Jerrum [12] combined the restriction that H is a tree with the restriction that counting is modulo 2. Their result will be important for our work, so we next give the definitions that we need to state their result.

The complexity class $\oplus$P [33, 22] contains all problems of the form "compute $f(x) \bmod 2$" such that computing $f(x)$ is a problem in #P. Toda [35] showed that there is a randomised polynomial-time reduction from every problem in the polynomial hierarchy to some problem in $\oplus$P. Thus, $\oplus$P-hardness is viewed as a stronger kind of intractability than NP-hardness. We use $\oplus\textsc{Hom}(H)$ to denote the computational problem of computing the number of homomorphisms from G to H, modulo 2, given an input graph G. It is immediate from the definition that $\oplus\textsc{Hom}(H)$ is in $\oplus$P.

The *involution-free reduction* of a graph H, from [12], is defined as follows. An *involution* σ of H is an automorphism of H whose order is at most 2 (that is, $\sigma \circ \sigma$ is the identity permutation). An involution is *non-trivial* if it is not the identity permutation. A graph H is *involution-free* if it has no non-trivial involutions. H^σ denotes the subgraph of H induced by the fixed points of σ (the vertices v with $\sigma(v) = v$). We write $H \to K$ if there is a non-trivial involution σ of H such that $K = H^\sigma$. The relation $\to^*$ is the reflexive-transitive closure of the relation $\to$. Thus, $H \to^* K$ means that either $K = H$, or there is a positive integer j and a sequence $H_1, \ldots, H_j$

and counting homomorphisms with restrictions such as surjectivity. These generalisations and restrictions are beyond the scope of this paper.

of graphs such that $H = H_1$, $K = H_j$ and, for all $i \in [j]$, $H_i \to H_{i+1}$. Faben and Jerrum [12, Theorem 3.7] showed that every graph H has, up to isomorphism, exactly one involution-free graph H^* such that $H \to^* H^*$. This graph H^* (labelled in a canonical way) is the *involution-free reduction* of H. The relevance of the involution-free reduction is given by the following theorem.

THEOREM 1.1 ([12, Theorem 3.4]). *For all graphs G and H, the number of homomorphisms from G to H has the same parity as the number of homomorphisms from G to H^*.*

Thus, the computational problem $\oplus\textsc{Hom}(H)$ reduces to $\oplus\textsc{Hom}(H^*)$. Faben and Jerrum made the following conjecture [12].

CONJECTURE 1.2 (Faben-Jerrum Conjecture). *Let H be a graph. If its involution-free reduction H^* has at most one vertex, then $\oplus\textsc{Hom}(H)$ can be solved in polynomial time. Otherwise, $\oplus\textsc{Hom}(H)$ is $\oplus$P-complete.*

The following progress has been made on the Faben-Jerrum conjecture.

- Faben and Jerrum [12, Theorem 3.8, Theorem 6.1] proved the conjecture for the case where every connected component of H is a tree.
- Göbel, Goldberg and Richerby [18, Theorem 3.8] proved the conjecture for the case where every connected component of H is a cactus graph, which is a connected, simple graph in which every edge belongs to at most one cycle.
- Göbel, Goldberg and Richerby [19, Theorem 1.2] proved the conjecture for the case where H is a simple graph whose involution-free reduction H^* is square-free (meaning that it has no 4-cycle).

The cactus-graph result generalises the tree result, and is incomparable with the square-free result.

1.2. Contributions and Techniques. Our first (and main) contribution is to prove the Faben-Jerrum conjecture for every simple graph H that does not have a K_4-minor.

Here, K_4 denotes the complete graph with four vertices. The concept of graph minors is well known (see, for example, [7]). In short, a graph H is K_4-minor-free if K_4 cannot be obtained from H by a sequence of vertex deletions, edge deletions, and edge contractions (removing any self-loops and multiple-edges that are formed by the contraction). Graph classes based on excluded minors form the basis of the graph structure theory of Robertson and Seymour (see [32]).

The class of K_4-minor-free graphs is a rich and well-studied class. It is equivalent to the class of graphs with treewidth at most 2 and it includes all outerplanar and series-parallel graphs [8].

Both trees and cactus graphs are K_4-minor free, so our result subsumes the tree result of Faben and Jerrum [12] and also the cactus-graph result of Göbel et al. [18]. K_4-minor-free graphs can contain a 4-cycle and, going the other way, square-free graphs can have a K_4-minor. Thus, our result is incomparable with the result of [19]. (As a more minor contribution, our techniques also give a shorter proof of the result of [19] — see Remark 5.8.)

Our second contribution is to extend $\oplus$P-hardness, using the randomised version of the Exponential Time Hypothesis of Impagliazzo and Paturi (rETH) to rule out subexponential algorithms. In order to state our result, we first state the hypothesis.

CONJECTURE 1.3 (rETH, [27]). *There is a positive constant c such that no algorithm, deterministic or randomised, can decide the satisfiability of an n-variable*

3-SAT *instance in time* $\exp(c \cdot n)$.

Using the rETH, we can now state our main result. Here (and in the rest of the paper) we denote the size of the input graph G as $|G| = |V(G)| + |E(G)|$.

THEOREM 1.4. *Let H be a simple graph whose involution-free reduction H^* is K_4-minor free. If H^* contains at most one vertex, then $\oplus\textsc{Hom}(H)$ can be solved in polynomial time. Otherwise, $\oplus\textsc{Hom}(H)$ is $\oplus$P-complete and, assuming the randomised Exponential Time Hypothesis, it cannot be solved in time* $\exp(o(|G|))$.

As an example of an application of Theorem 1.4, consider the following K_4-minor-free graphs H_1 and H_2.

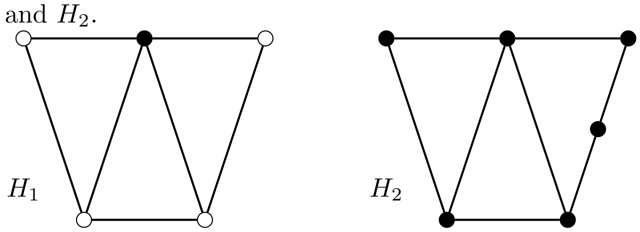

The graph H_1 has a non-trivial involution whose only fixed-point is the solid vertex, so H_1^* has one vertex. By Theorem 1.4, $\oplus\textsc{Hom}(H_1)$ can be solved in polynomial time. The graph H_2 does not have any non-trivial involutions, so $H_2^* = H_2$. By Theorem 1.4, $\oplus\textsc{Hom}(H_2)$ is $\oplus$P-complete and it cannot be solved in time $\exp(o(|G|))$, unless the rETH fails.

Before describing our techniques, we mention that they lead easily to a couple of other results — a proof of the Faben-Jerrum conjecture for graphs whose involution-free reduction have degree at most 3 (Theorem 10.2) and a complete complexity classification for counting list homomorphisms modulo 2 (Theorem 11.3).

Technical Overview. Given Theorem 1.1, we focus on the case where H is involution-free. In general, our proof proceeds in two steps. Given an involution-free K_4-minor-free graph H, in step 1 we try to find a biconnected component of H, let us call it B, that allows us to derive $\oplus$P-hardness of $\oplus\textsc{Hom}(H)$ by exploiting the *local* structure of B to construct a reduction from counting independent sets, modulo 2. The latter problem, denoted by $\oplus$IS, is known to be $\oplus$P-complete [36] and cannot be solved in subexponential time, unless the rETH fails [5].

A careful analysis of biconnected and K_4-minor-free graphs, which crucially relies on the absence of non-trivial involutions, shows that the first step is always possible, unless all biconnected components of H have a very restricted form; examples are depicted in Figure 1.

The second step of the proof exploits the global structure of H and deals with the case where step 1 fails. Note that all of the depicted biconnected components have non-trivial involutions; consider for example the involution given by swapping the vertices x and y in Figure 1. Since the overall graph H is promised to be free of such involutions, we infer that at least one of x and y has a neighbour in a further biconnected component of H, which will allow us to successively construct a global "walk-like" structure in H that eventually yields a reduction from $\oplus$IS.

We consider the construction of those global substructures as our main technical contribution. While the formal specifications are beyond the scope of the introduction, we give an illustrated example which we hope gives some flavour of the graph theory that we will encounter in this work:

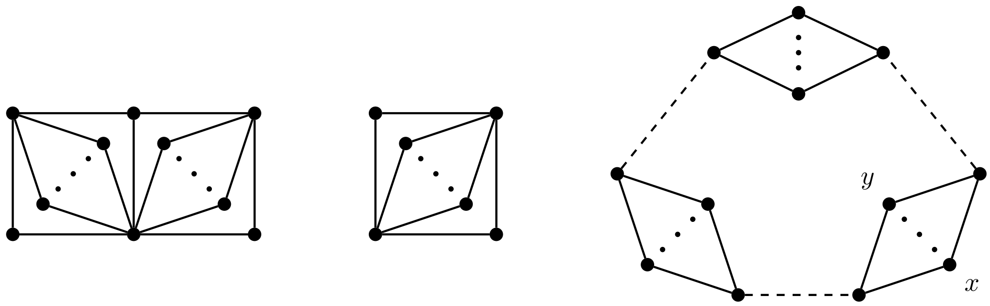

Fig. 1: Examples of three types of biconnected and K_4-minor-free graphs that, if viewed as biconnected components, do not yield a reduction from ⊕IS. We will re-encounter those graphs as "impasses" *(left)*, "diamonds" *(centre)*, and "obstructions" *(right)*.

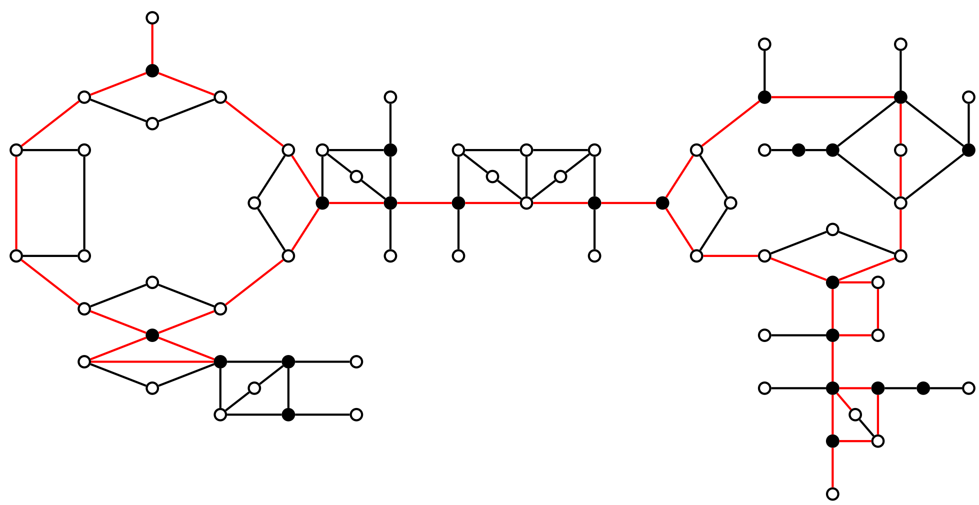

The above illustration depicts a K_4-minor-free graph H' without non-trivial involutions, together with a subgraph, highlighted in red, that allows for a reduction from ⊕IS. Solid vertices depict articulation points, i.e., vertices that lie in the intersection of at least 2 biconnected components. Note that each biconnected component of H' that is not an edge is of one of the three types given in Figure 1. Also, each biconnected component of H' has an involution. These non-trivial involutions prevent us from exploiting the local structure of the biconnected components to derive ⊕P-hardness. Instead, we will see that the highlighted subgraph is what makes $\oplus\textsc{Hom}(H')$ hard.

In the next section we provide an overview of the general framework that allows us to reduce ⊕IS to $\oplus\textsc{Hom}(H)$. The structures used in such reductions are captured by the so-called hardness gadgets introduced by Göbel, Goldberg and Richerby [18, 19]. Prior applications of hardness gadgets could only be used to construct a reduction from ⊕IS to $\oplus\textsc{Hom}(H)$ if H has certain local substructures, based around a path or a cycle. In contrast, our analysis will establish global walks such as the one highlighted in H'. As far as we can tell, none of the prior machinery [12, 18, 19, 29] is capable of proving the ⊕P-hardness of $\oplus\textsc{Hom}(H')$, however, this will follow as a result of our abstract consideration of global substructures of K_4-minor-free graphs.

2. Warm-up: useful Ideas from Previous Papers — Retractions and Hardness Gadgets . Instead of directly reducing $\oplus$IS to $\oplus\textsc{Hom}(H)$, it is useful to consider the intermediate problem $\oplus\textsc{Ret}(H)$, the problem of counting *retractions*[2] to H, modulo 2. Given a graph H, a *partially H-labelled graph* $J = (G, \tau)$ consists of an *underlying graph* G and a corresponding *pinning function* τ, which is a partial function from $V(G)$ to $V(H)$. A homomorphism from J to H is a homomorphism h from G to H such that, for all vertices v in the domain of τ, $h(v) = \tau(v)$.

A homomorphism from a partially H-labelled graph J to H is also called a *retraction* to H because we can think of the pinning function τ as a way of identifying an induced subgraph H of G which must "retract" to H under the action of the homomorphism — see [13] for details. We use $\oplus\textsc{Ret}(H)$ to denote the computational problem of computing the number of homomorphisms from J to H, modulo 2, given as input a partially H-labelled graph J.

It is known [19] that $\oplus\textsc{Ret}(H)$ reduces to $\oplus\textsc{Hom}(H)$ whenever H is involution-free. Since τ allows us to pin vertices of G to vertices of H arbitrarily, it is much easier to construct a reduction from $\oplus$IS to $\oplus\textsc{Ret}(H)$ than to construct a direct reduction from $\oplus$IS to $\oplus\textsc{Hom}(H)$.

Consider the following example. Suppose that H is the 4-vertex path (o, s, i, x) and that our goal is to reduce $\oplus$IS to $\oplus\textsc{Ret}(H)$. Let G be an input to $\oplus$IS. That is, G is a graph whose independent sets we wish to count, modulo 2. For ease of presentation, suppose that G is bipartite,[3] that is, the vertices of G can be partitioned into two independent sets U and V. Let $\widehat{G}$ be the graph obtained from G by adding two additional vertices u and v, and by connecting u to all vertices in U, and v to all vertices in V, respectively. Let τ be the pinning function defined by $\tau(u) = s$ and $\tau(v) = i$. We provide an illustration of the construction in Figure 2.

Observe that any homomorphism φ from $(\widehat{G}, \tau)$ to H must map every vertex in U to either o or i, and every vertex in V to either s or x. Since H has no edge from o to x, the definition of homomorphism ensures that $\varphi^{-1}(o) \cup \varphi^{-1}(x)$ is an independent set of G. It is easy to verify that the function $\varphi \mapsto \varphi^{-1}(o) \cup \varphi^{-1}(x)$ is a bijection between the homomorphisms from $(\widehat{G}, \tau)$ to H and the independent sets of G, which gives a reduction from (bipartite) $\oplus$IS to $\oplus\textsc{Ret}(H)$.

The observant reader might notice that the 4-vertex path has a non-trivial involution, and thus, we cannot further reduce $\oplus\textsc{Ret}(H)$ to $\oplus\textsc{Hom}(H)$ in this case.[4] However, the construction works for *any* graph H with an induced path (o, s, i, x) such that s and i each only have two neighbours.

The notion of a *hardness gadget*, which we formally introduce in Section 4, is essentially a generalisation of the previous construction. For example, we could substitute each of o, s, i and x with an odd number of copies, since we are only interested in the parity of the number of independent sets. Furthermore, we could identify o and x, since we only need the edge $\{o, x\}$ to be absent in H. A more sophisticated generalisation is obtained by observing that we can, to some extent, substitute the edges $\{o, s\}$, $\{s, i\}$ and $\{i, x\}$ with more complicated graphs, e.g. with length-2 paths, if we substitute the edges in $\widehat{G}$ accordingly. Finally, observe that the construction

[2] In some definition versions a retraction is surjective. However, for algorithmic problems this surjectivity requirement is not important [13, 15]

[3] The case of general graphs will be discussed later in the paper.

[4] In fact, the problem $\oplus\textsc{Hom}(H)$ is trivial when H is the 4-vertex path since the number of homomorphisms will always be even.

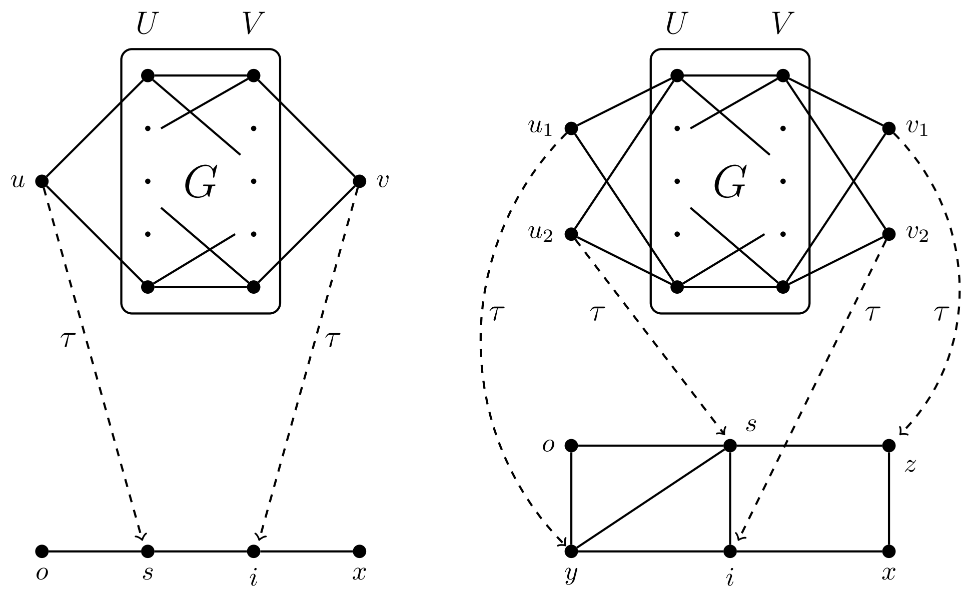

Fig. 2: Illustration of the reduction from (bipartite) $\oplus$IS to $\oplus$RET(H) where H is the 4-vertex path *(left)*, and H is the graph H_2 from page 4 *(right)*.

$(\widehat{G}, \tau)$ uses the partial function τ in a very simple manner: By adding a common neighbour u for all vertices in U and setting $\tau(u) = s$, the construction enforces the constraint that any homomorphism from $(\widehat{G}, \tau)$ to H must map every vertex in U to a neighbour of s. More sophisticated constructions will allow us to enforce much stronger constraints on homomorphisms. We will need this flexibility to construct reductions from $\oplus$IS to $\oplus$RET(H) for more general graphs H.

We conclude by making a generalisation explicit for one further example — the graph H_2 from page 4. We provide a more convenient drawing of H_2, including a labelling of its vertices and an illustration of the reduction in Figure 2. Again, we will assume for ease of presentation that the input G to $\oplus$IS is bipartite. To construct $\widehat{G}$, we add two additional vertices u_1 and u_2 and make them adjacent to every vertex in U. Similarly, we add two additional vertices v_1 and v_2 and make them adjacent to every vertex in V. Let τ be the pinning function defined by $\tau(u_1) = y$, $\tau(u_2) = s$, $\tau(v_1) = z$, and $\tau(v_2) = i$.

Consider any homomorphism φ from $(\widehat{G}, \tau)$ to H_2. Since φ is edge-preserving, it must map every vertex in U to a common neighbour of s and y in H_2. Consequently, $\varphi(U) \subseteq \{o, i\}$. Similarly, we obtain $\varphi(V) \subseteq \{s, x\}$. Again, it is easy to see that the mapping $\varphi \mapsto \varphi^{-1}(o) \cup \varphi^{-1}(x)$ is a bijection between the homomorphisms from $(\widehat{G}, \tau)$ to H and the independent sets of G, which gives a reduction from (bipartite) $\oplus$IS to $\oplus$RET(H).

Note that the second example, while being less straightforward than the first, is still a very simple reduction. The proof of Theorem 1.4 requires us to consider much more intricate "hardness gadgets"; the necessary tools will be carefully introduced in Sections 5 and 6.

3. Proof Outline and Organisation of the Paper. We start with the formal definitions that we need in Section 4; in particular we set up the framework of hardness gadgets. Section 5, our "toolbox", presents the most important class of hardness gadgets that we use.

Sections 6-9 constitute the proof of our main result and should be considered the technical core of this paper. In Section 6 we deal with biconnected K_4-minor-free graphs that are additionally chordal bipartite graphs (that is, they have the property that every induced cycle is a square). The reason for our separate treatment of these graphs is that our main gadget from Section 5 cannot be applied to graphs without an induced cycle of length $\neq 4$. We identify two families of such graphs, *impasses* and *diamonds*, that prevent us from constructing a local hardness gadget; examples of an impasse and a diamond are depicted in Figure 1.

After that, we dedicate Section 7 to the analysis of K_4-minor-free graphs that contain certain sequences of biconnected components, each of which is either an edge, an impasse, or a diamond. In Section 8 we consider biconnected K_4-minor-free graphs that are not necessarily chordal bipartite. We identify another family of graphs that does not allow for a local, i.e., an "internal", hardness gadget; we call such graphs *obstructions*; obstructions always contain an induced cycle of length other than 4, and an example of an obstruction is depicted in Figure 1.

In combination, Sections 7 and 8 reveal the structure of involution-free K_4-minor-free graphs that do not allow for a local hardness gadget. In Section 9 we use this structure, which allows us to constructively prove the existence of *global* hardness gadgets for all remaining K_4-minor-free graphs without non-trivial involutions. Our main theorem, Theorem 1.4, follows.

In Sections 10 and 11 we explore the applicability of our machinery to further classes of graphs and problems: Section 10 presents a full classification for counting homomorphisms to graphs of degree at most 3, modulo 2. Section 11 considers the problem of counting list homomorphisms, modulo 2, a variation of the homomorphism problem that generalises retractions as follows: Let H be a fixed graph. The problem $\oplus\text{LH}\textsc{om}(H)$ expects as input a graph G and a function τ that maps every vertex of G to a list of vertices of H. The goal is then to compute the parity of the number of homomorphisms from G to H which additionally map every vertex v of G to a vertex contained in $\tau(v)$. We provide a full classification of $\oplus\text{LH}\textsc{om}(H)$ for all graphs H, even if self-loops are allowed.

Finally, in Section 12, we provide an index containing the most important symbols and definitions.

4. Preliminaries. An index of notation and terminology is in Section 12. Given a positive integer q let $[q] = \{1, \dots, q\}$. Given a finite set S, we write $|S|$ for its cardinality.

Graph theory. Graphs in this work are simple, that is, without multiple edges, and do not contain self-loops, unless stated otherwise. The size of a graph G is defined as $|G| = |V(G)| + |E(G)|$. Given a graph H and a subset S of its vertices, we write $H[S]$ for the subgraph of H *induced* by S.

Given a non-negative integer k, a *walk* of length k in a graph H is a sequence of (not necessarily distinct) vertices $(v_0, \dots, v_k)$ such that, for all $i \in [k]$, $\{v_{i-1}, v_i\} \in E(H)$. The walk is *closed* if $v_0 = v_k$. Note that for $k = 0$, the single vertex (v_0) is a closed walk of length 0. A *path* of length k is a walk of length k for which $v_0, \dots, v_k$ are distinct. For $k \geq 3$, a *cycle* of length k is a closed walk of length k such that $v_1, \dots, v_k$ are distinct. A *square* is a cycle of length 4.

For $i, j \in \{0, \ldots, k\}$ with $i \le j$, we say that $(v_i, v_{i+1} \ldots, v_j)$ is a *subwalk* of $(v_0, \ldots, v_k)$. For vertices $u, v \in V(H)$, $\mathsf{dist}_H(u, v)$ is the length of a shortest path between u and v.

Definition 4.1 (chordal bipartite graph, see e.g. [23]). A graph in which every induced cycle is a square is called a *chordal bipartite graph.*

Given a graph H and a vertex $v \in V(H)$, we write $\Gamma_H(v)$ for the *neighbourhood* of v in H and we write $\deg_H(v)$ for the *degree* of v. That is, $\Gamma_H(v) = \{u \in V(H) \mid \{u, v\} \in E(H)\}$ and $\deg_H(v) = |\Gamma_H(v)|$. Given a subset S of $V(H)$, we set $\Gamma_H(S) = \bigcap_{v \in U} \Gamma_H(v)$ and note that $\Gamma_H(v) = \Gamma_H(\{v\})$.

Definition 4.2 (walk-neighbour-set). Let $W = (w_0, \ldots, w_{q-1}, w_0)$ be a closed walk in a graph H. We use $N_{W,H}(w_i)$ to denote $\Gamma_H(w_{i-1}) \cap \Gamma_H(w_{i+1})$, where the indices are taken modulo q. We refer to the sets $N_{W,H}(w_0), \ldots, N_{W,H}(w_{q-1})$ as the *walk-neighbour-sets* of W in H.

Definition 4.3 (articulation point, biconnected, block-cut tree). An *articulation point* of a graph is a vertex whose removal increases the number of connected components. A graph is *biconnected* if it has at least 2 vertices and has no articulation point. A *biconnected component* is a maximal biconnected subgraph.

Let H be a connected graph. The *block-cut tree* of H is the tree $\mathsf{BC}(H)$ that has a vertex for each biconnected component of H (such vertices are called *blocks*) and a vertex for each articulation point of H (such vertices are also called *cut vertices*) such that T has an edge between each biconnected component B and each articulation point a in B.

Partially labelled graphs. Let H be a graph. Recall from Section 2 that a partially H-labelled graph $J = (G, \tau)$ consists of an underlying graph G and a corresponding pinning function τ, which is a partial function from $V(G)$ to $V(H)$. A vertex v in the domain of the pinning function is said to be *pinned*, *pinned to* $\tau(v)$, or a $\tau(v)$*-pin*. We write partial functions as sets of pairs, for example, writing $\tau = \{a \mapsto s, b \mapsto t\}$ for the partial function τ with domain $\{a, b\}$ such that a is an s-pin and b is a t-pin. The size of a partially H-labelled graph $J = (G, \tau)$ is defined as $|J| = |G|$.

Homomorphisms and Counting (mod 2). Given graphs G and H, $\hom(G \to H)$ denotes the *set of homomorphisms* from G to H and $\hom(J \to H)$ denotes the set of homomorphisms from J to H.

It will sometimes be convenient to consider a graph G together with some number of distinguished vertices $x_1, \ldots, x_r$ of G. We denote such a graph by $(G, x_1, \ldots, x_r)$. The distinguished vertices need not be distinct. A homomorphism from a graph $(G, x_1, \ldots, x_r)$ to $(H, y_1, \ldots, y_r)$ is a homomorphism h from G to H with the property that, for each $i \in [r]$, $h(x_i) = y_i$. We write $\hom((G, x_1, \ldots, x_r) \to (H, y_1, \ldots, y_r))$ for the set of such homomorphisms.

Given a partially labelled graph $J = (G, \tau)$ and distinguished vertices $x_1, \ldots, x_r$ of G that are not in the domain of τ, a homomorphism from $(J, x_1, \ldots, x_r)$ to $(H, y_1, \ldots, y_r)$ is a homomorphism from $J' = (G, \tau \cup \{x_1 \mapsto y_1, \ldots, x_r \mapsto y_r\})$ to H. The set of such homomorphisms is denoted by $\hom((J, x_1, \ldots, x_r) \to (H, y_1, \ldots, y_r))$.

Useful tools. The following theorem of Göbel, Goldberg and Richerby will be of crucial importance in this work, as it will allow us to derive hardness of $\oplus\textsc{Hom}(H)$ from hardness of $\oplus\textsc{Ret}(H)$.

Theorem 4.4 ([19, Theorem 3.1]). *Let H be an involution-free graph. Then there is an algorithm with oracle access to $\oplus\textsc{Hom}(H)$ that takes as input a partially H-labelled graph J and computes $|\hom(J \to H)| \mod 2$ in time $\mathrm{poly}(|J|)$. The size* *of the input to every oracle query is $O(|J|)$.*

The statement of Theorem 4.4 in [19, Theorem 3.1] does not mention the fact that the size of the input to every oracle query is $O(|J|)$. Nevertheless, it is easy to see, by examining the proof in [19] that this linearity requirement is met (without making any changes to the proof). The reason that we introduce this linearity constraint is so that our hardness results can also rule out subexponential-time algorithms for $\oplus\textsc{Hom}(H)$ in the $\oplus$P-hard cases, using the rETH.

The following theorem of Faben and Jerrum will also be useful, as it will allow us to focus on connected graphs. The statement of [12, Theorem 6.1] does not mention the linearity requirement on the size of oracle queries, but this requirement does not present any difficulties. Faben and Jerrum's proof is given in a slightly different setting (pinning to orbits of vertices of H rather than to vertices) so, for completeness, we give a short proof.

Lemma 4.5 ([12, Theorem 6.1]). *Let H be an involution-free graph and let H' be a connected component of H. Then there exists an algorithm with oracle access to $\oplus\textsc{Hom}(H)$ that takes as input a graph G and computes $|\mathrm{hom}(G \to H')| \mod 2$ in time $\mathrm{poly}(|G|)$. The size of every oracle query is $O(|G|)$.*

Proof. Let G be a graph. If G is the empty graph then the algorithm returns 1, which is the number of homomorphisms from G to H'. Otherwise, there exists a vertex $u \in V(G)$. For each $v \in V(H')$ we define the partially H'-labelled $J_v = (G, \{u \mapsto v\})$. Note that $|\mathrm{hom}(G \to H')| = \sum_{v \in V(H')} |\mathrm{hom}(J_v \to H)|$.

By Theorem 4.4, there is an algorithm A with oracle access to $\oplus\textsc{Hom}(H)$ that takes as input a partially H-labelled graph J and computes $|\mathrm{hom}(J \to H)| \mod 2$ in time $\mathrm{poly}(|J|)$ such that the size of every oracle query is bounded by $O(|J|)$. Our algorithm uses algorithm A as a subroutine to compute the parity of $|\mathrm{hom}(J_v \to H)|$ for each $v \in V(H')$. This requires $|V(H')|$ executions of the subroutine A. Thus, the algorithm runs in time

$$O\Big(\sum_{v \in V(H')} \mathrm{poly}(|J_v|)\Big) = \mathrm{poly}(|G|).$$

Moreover, for each $v \in V(H')$, the size of each $\oplus\textsc{Hom}(H)$ oracle query is bounded by $O(|J_v|) = O(|G|)$. ∎

Hardness Gadgets. The following is a slightly generalised version of the *hardness gadget* introduced in [19, Definition 4.1]. The only difference between their definition and ours is that they require the sets I and S to have size 1.

Definition 4.6. [19, Definition 4.1] A *hardness gadget* for a graph H is a tuple $(I, S, (J_1, y), (J_2, z), (J_3, y, z))$ that consists of odd-cardinality sets $I, S \subseteq V(H)$ together with three connected, partially H-labelled graphs with distinguished vertices (J_1, y), (J_2, z) and (J_3, y, z) that satisfy certain properties as explained below. Let

$$\begin{aligned}
\Omega_y &= \{a \in V(H) \mid |\mathrm{hom}((J_1, y) \to (H, a))| \text{ is odd}\},\\
\Omega_z &= \{b \in V(H) \mid |\mathrm{hom}((J_2, z) \to (H, b))| \text{ is odd}\}, \text{ and}\\
\Sigma_{a,b} &= \mathrm{hom}((J_3, y, z) \to (H, a, b)).
\end{aligned}$$

The properties that we require are the following.

1. $|\Omega_y|$ is even and $I \subset \Omega_y$.
2. $|\Omega_z|$ is even and $S \subset \Omega_z$.

3. For each $i \in I$, $o \in \Omega_y \setminus I$, $s \in S$ and each $x \in \Omega_z \setminus S$,
   - $|\Sigma_{o,x}|$ is even.
   - $|\Sigma_{i,s}|$, $|\Sigma_{o,s}|$ and $|\Sigma_{i,x}|$ are odd.

The following theorem of Göbel, Goldberg and Richerby establishes intractability of $\oplus\textsc{Ret}(H)$ whenever H has a hardness gadget.

THEOREM 4.7 ([19, Theorem 4.2]). *Let H be an involution-free graph that has a hardness gadget. Then $\oplus\textsc{Ret}(H)$ is $\oplus$P-hard. Also, assuming the randomised Exponential Time Hypothesis, $\oplus\textsc{Ret}(H)$ cannot be solved in time* $\exp(o(|J|))$.

Proof. Although the hardness gadgets from [19] are more constrained than the ones that we use, the proof of [19, Theorem 4.2] establishes the $\oplus$P-hardness in Theorem 4.7 with only very minor changes, which we now describe.

As noted in the introduction, Valiant [36] showed that the problem $\oplus$IS is $\oplus$P-complete. The proof of [19, Theorem 4.2] gives a polynomial-time Turing reduction from $\oplus$IS to $\oplus\textsc{Ret}(H)$. The reduction uses G and the hypothesised hardness gadget for H to construct a partially H-labelled graph J such that the number of independent sets of G, which we denote $|\mathcal{I}(G)|$, is equal, modulo 2, to $|\hom(J \to H)|$. The reduction concludes by making a single oracle call to $\oplus\textsc{Ret}(H)$ with input J.

In our case, the construction of J is exactly as it is in [19]. The proof that $|\mathcal{I}(G)| = |\hom(J \to H)| \mod 2$ needs only a very minor modification to account for the fact that the sets I and S in the hardness gadget may have more than one element. At some point in the proof of [19], it is argued that a certain quantity $n(a, a')$ is even if a and a' are both in I, and odd otherwise. This is still true even when I and S contain more than one element — it follows from item 3 in the definition of hardness gadget (and from the fact that I and S have odd cardinality).

The final sentence in the statement of Theorem 4.7, asserting that $\oplus\textsc{Ret}(H)$ cannot be solved in time $\exp(o(|J|))$ unless the rETH fails, was not contained in the original theorem of [19], however it follows immediately from the fact that $|J| = O(|G|)$ (which is easily checked) and from the fact that $\oplus$IS cannot be solved in time $\exp(o(|G|))$, unless the rETH fails, which was proved by Dell, Husfeldt, Marx, Taslaman and Wahlen [5]. [5]

More precisely, for establishing the conditional lower bound for $\oplus\textsc{Ret}(H)$, let us assume for contradiction that we can solve $\oplus\textsc{Ret}(H)$ in time $\exp(o(|J|))$. We obtain an algorithm for $\oplus$IS that, on input G, runs the (polynomial-time) Turing-reduction from [19, Theorem 4.2] and then simulates each oracle query J for $\oplus\textsc{Ret}(H)$ in time $\exp(o(|J|))$; note that this simulation is possible by our assumption. Since each oracle query J has size at most $O(|G|)$, the total running time of our algorithm for $\oplus$IS is bounded by $\mathsf{poly}(|G|) \cdot \exp(o(|G|)) = \exp(o(|G|))$, contradicting rETH. This concludes the proof of the conditional lower bound. ∎

5. Toolbox.

5.1. Path Gadget.

We will use the following path gadget, which is called a "caterpillar gadget" in [19].

Definition 5.1. Given a path $P = (v_0, \ldots, v_q)$ in H with $q \geq 1$, define the *path gadget* $J_P = (G, \tau)$ as follows. $V(G) = \{u_1, \ldots, u_{q-1}, w_1, \ldots, w_{q-1}, y, z\}$ and G is

[5] In more detail, Dell et al. established that counting independent sets cannot be done in time $\exp(o(|E(G)|))$, unless the rETH fails [5, Theorem 1.2]. They point out explicitly that their reduction also works in the case of counting modulo 2. Furthermore, their reduction always yields a graph without isolated vertices — for such graphs we have $|E(G)| = \Theta(|G|)$.

the path $(y, u_1, \ldots, u_{q-1}, z)$ together with edges $\{u_j, w_j\}$ for $j \in [q-1]$. $\tau = \{w_1 \mapsto v_1, \ldots, w_{q-1} \mapsto v_{q-1}\}$.

We will use the following lemma of Göbel, Goldberg and Richerby. The original lemma was stated for square-free graphs, but the proof only uses the fact that no edge of P is part of a square in H.

Lemma 5.2 ([19, Lemma 4.5]). *For an integer $q \geq 1$, let $P = (v_0, \ldots, v_q)$ be a path in a graph H. Suppose that no edge of P is part of a square in H and that $\deg_H(v_j)$ is odd for all $j \in [q-1]$. Let $\Omega_y \subseteq \Gamma_H(v_0)$ and $\Omega_z \subseteq \Gamma_H(v_q)$, with $I = \{v_1\} \subset \Omega_y$ and $S = \{v_{q-1}\} \subset \Omega_z$. For $i = v_1$, $s = v_{q-1}$ and for each $o \in \Omega_y \setminus I$ and $x \in \Omega_z \setminus S$ we have the following:*

- $|\hom((J_p, y, z) \to (H, o, x))| = 0$,
- $|\hom((J_p, y, z) \to (H, o, s))| = 1$,
- $|\hom((J_p, y, z) \to (H, i, x))| = 1$, *and*
- $|\hom((J_p, y, z) \to (H, i, s))|$ *is odd.*

5.2. Cycle Gadget. We will use the following cycle gadget, which is a generalisation of the cycle gadget in [29].

Definition 5.3 (Cycle gadget). For an integer $q \geq 3$, let $\mathcal{C} = (\mathcal{C}_0, \ldots, \mathcal{C}_{q-1})$ where, for $i = 0, \ldots, q-1$, s_i is a positive integer and $\mathcal{C}_i = \{c_i^1, \ldots, c_i^{s_i}\}$ is a set of s_i vertices. We define the *cycle gadget* $J_{\mathcal{C}} = (G, \tau)$ as follows (see Figure 3). For $i = 0, \ldots, q-1$, let $U_i = \{u_i^1, \ldots, u_i^{s_i}\}$ be a set of s_i vertices. Then $V(G) = \{v_0, \ldots, v_{q-1}\} \cup U_0 \cup \cdots \cup U_{q-1}$ (where all named vertices are assumed to be distinct) and $E(G) = \{\{v_i, v_{i+1 \bmod q}\} \mid i \in \{0, \ldots, q-1\}\} \cup \{\{v_i, u_i^j\} \mid i \in \{0, \ldots, q-1\}, j \in \{1, \ldots, s_i\}\}$. $\tau = \{u_i^j \mapsto c_i^j \mid \forall i \in \{0, \ldots q-1\}, j \in \{1, \ldots s_i\}\}$.

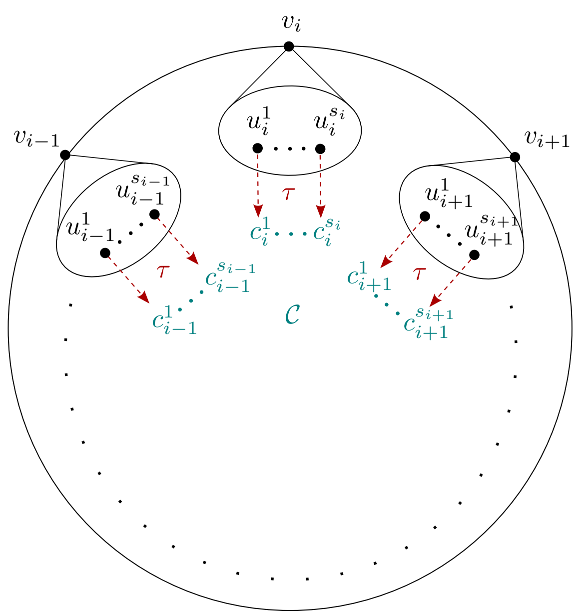

Fig. 3: The cycle gadget $J_{\mathcal{C}}$ is depicted in black. The corresponding pinning function τ is depicted in red, it maps to the vertices of $\mathcal{C}$, depicted in blue.

In fact, we will also need a further generalisation of the cycle gadget from Definition 5.3.

Definition 5.4 (Generalised cycle gadget). Intuitively, we generalise the cycle gadget by attaching, at each vertex v_i, a gadget J_i. Let H be a graph. For an integer $q \geq 3$, let $\mathcal{C} = (\mathcal{C}_0, \ldots, \mathcal{C}_{q-1})$ where, for $i = 0, \ldots, q-1$, s_i is a positive integer and $\mathcal{C}_i = \{c_i^1, \ldots, c_i^{s_i}\}$ is a set of s_i vertices of H. Let $J_{\mathcal{C}}$ be the cycle gadget from Definition 5.3. Let $(J_0, z_0), \ldots, (J_{q-1}, z_{q-1})$ be partially H-labelled graphs with distinguished vertices. Then the *generalised cycle gadget* $J(J_{\mathcal{C}}, J_0, \ldots, J_{q-1})$ is the gadget obtained from $J_{\mathcal{C}}, J_0, \ldots, J_{q-1}$ by identifying, for each $i \in \{0, \ldots, q-1\}$ the vertex v_i from $J_{\mathcal{C}}$ with the vertex z_i from J_i.

Lemma 5.5. *For an integer $q \geq 3$, let H be a graph which contains sets of vertices $\mathcal{C}_0, \ldots, \mathcal{C}_{q-1}$ (not necessarily disjoint or even distinct). Let $(J_0, z_0), \ldots, (J_{q-1}, z_{q-1})$ be partially H-labelled graphs with distinguished vertices, and, for each $i \in \{0, \ldots, q-1\}$, let $\Omega_i = \{a \in V(H) \mid |\hom((J_i, z_i) \to (H, a))| \text{ is odd}\}$. Suppose that for all $i \in \{0, \ldots, q-1\}$ we have the following.*

(L5.5.1) $|\mathcal{C}_{i-1 \bmod q} \cap \Omega_i|$ and $|\mathcal{C}_{i+1 \bmod q} \cap \Omega_i|$ are odd.

(L5.5.2) If $w \in \mathcal{C}_{i-1 \bmod q}$ then $\Gamma_H(w) \cap \Gamma_H(\mathcal{C}_{i+1 \bmod q}) = \mathcal{C}_i$.

(L5.5.3) If $w \in \mathcal{C}_{i+1 \bmod q}$ then $\Gamma_H(\mathcal{C}_{i-1 \bmod q}) \cap \Gamma_H(w) = \mathcal{C}_i$.

(L5.5.4) There is no walk of the form $D = (d_0, \ldots, d_{q-1}, d_0)$ such that, for all $i \in \{0, \ldots, q-1\}$, $d_i \in \Gamma_H(\mathcal{C}_i) \setminus (\mathcal{C}_{i-1 \bmod q} \cup \mathcal{C}_{i+1 \bmod q})$.

Let $J_{\mathcal{C}}$ be the cycle gadget (Definition 5.3) and let $J^* = J(J_{\mathcal{C}}, J_0, \ldots, J_{q-1})$ be the generalised cycle gadget (Definition 5.4) Then, for all $k \in \{0, \ldots, q-1\}$,

$$\{a \in V(H) \mid |\hom((J^*, v_k) \to (H, a))| \text{ is odd}\} = (\mathcal{C}_{k-1 \bmod q} \cup \mathcal{C}_{k+1 \bmod q}) \cap \Omega_k.$$

Proof. To simplify notation, all indices in this proof are considered to be modulo q. For $a \in V(H)$, let $k \in \{0, \ldots, q-1\}$ and $h \in \hom((J^*, v_k) \to (H, a))$. By construction of J^* and the fact that h has to preserve edges, for all $i \in \{0, \ldots, q-1\}$, we obtain

- $h(v_i) \in \Gamma_H(\mathcal{C}_i)$,
- $h(v_i) \notin \mathcal{C}_i$ (since we do not allow self-loops in H),
- $h(v_i)$ is adjacent to $h(v_{i+1})$ in H,
- $h(v_i) \neq h(v_{i+1})$.

Consequently, it holds that $h(v_{i+1}) \in \Gamma_H(h(v_i)) \cap \Gamma_H(\mathcal{C}_{i+1})$. Suppose, for some $i \in \{0, \ldots, q-1\}$, that $h(v_i) \in \mathcal{C}_{i-1}$. Then, by (L5.5.2), we have $h(v_{i+1}) \in \mathcal{C}_i$. Therefore,

$$\text{If } h(v_i) \in \mathcal{C}_{i-1} \text{ then } h(v_{i+1}) \in \mathcal{C}_i. \tag{5.1}$$

Analogously, using (L5.5.3),

$$\text{If } h(v_i) \in \mathcal{C}_{i+1} \text{ then } h(v_{i-1}) \in \mathcal{C}_i. \tag{5.2}$$

Thus, if there exists some $\ell \in \{0, \ldots, q-1\}$ such that $h(v_\ell) \in \mathcal{C}_{\ell-1}$ then we can use (5.1) iteratively to obtain $h(v_i) \in \mathcal{C}_{i-1}$ for all $i \in \{0, \ldots, q-1\}$. In particular, $h(v_k) \in \mathcal{C}_{k-1}$. Analogously, if there exists some $\ell \in \{0, \ldots, q-1\}$ such that $h(v_\ell) \in \mathcal{C}_{\ell+1}$ then we can use (5.2) iteratively to obtain $h(v_i) \in \mathcal{C}_{i+1}$ for all $i \in \{0, \ldots, q-1\}$. This means that $h(v_k) \in \mathcal{C}_{k+1}$.

Suppose that $h(v_k) \notin \mathcal{C}_{k-1} \cup \mathcal{C}_{k+1}$. We have established that, using (5.1) and (5.2) iteratively, we obtain, for all $i \in \{0, \ldots, q-1\}$, $h(v_i) \notin \mathcal{C}_{i-1} \cup \mathcal{C}_{i+1}$ and consequently $h(v_i) \in \Gamma_H(\mathcal{C}_i) \setminus (\mathcal{C}_{i-1} \cup \mathcal{C}_{i+1})$. However, $(h(v_0), \ldots, h(v_{q-1}), h(v_0))$ is a walk in H, which gives a contradiction to (L5.5.4).

We have shown that $h(v_k) \in \mathcal{C}_{k-1} \cup \mathcal{C}_{k+1}$. Moreover, for each $a \in \mathcal{C}_{k-1}$, we have $|\hom((J^*, v_k) \to (H,a))| = |\hom((J_k, z_k) \to (H,a))| \cdot \prod_{i\in\{0,\ldots,q-1\}\setminus\{k\}} |\mathcal{C}_{i-1} \cap \Omega_i|$, which is odd if and only if $a \in \mathcal{C}_{k-1} \cap \Omega_k$ by (L5.5.1). The statement for $a \in \mathcal{C}_{k+1}$ is analogous. ∎

LEMMA 5.6. *For an integer $q \geq 3$, let H be a graph which contains sets of vertices $\mathcal{C}_0, \ldots, \mathcal{C}_{q-1}$ (not necessarily disjoint or even distinct). Let $(J_0, z_0), \ldots, (J_{q-1}, z_{q-1})$ be partially H-labelled graphs with distinguished vertices, and, for each $i \in \{0, \ldots, q-1\}$, let $\Omega_i = \{a \in V(H) \mid |\hom((J_i, z_i) \to (H,a))|$ is odd$\}$. Suppose that for all $i \in \{0, \ldots, q-1\}$ we have the following properties from the statement of Lemma 5.5.*

(L5.5.1) $|\mathcal{C}_{i-1 \bmod q} \cap \Omega_i|$ *and* $|\mathcal{C}_{i+1 \bmod q} \cap \Omega_i|$ *are odd.*

(L5.5.2) *If* $w \in \mathcal{C}_{i-1 \bmod q}$ *then* $\Gamma_H(w) \cap \Gamma_H(\mathcal{C}_{i+1 \bmod q}) = \mathcal{C}_i$.

(L5.5.3) *If* $w \in \mathcal{C}_{i+1 \bmod q}$ *then* $\Gamma_H(\mathcal{C}_{i-1 \bmod q}) \cap \Gamma_H(w) = \mathcal{C}_i$.

(L5.5.4) *There is no walk of the form* $D = (d_0, \ldots, d_{q-1}, d_0)$ *such that, for all* $i \in \{0, \ldots, q-1\}$, $d_i \in \Gamma_H(\mathcal{C}_i) \setminus (\mathcal{C}_{i-1 \bmod q} \cup \mathcal{C}_{i+1 \bmod q})$.

Furthermore, there exists $k \in \{0, \ldots, q-1\}$ such that

(L5.6.1) *there are no edges between* $\mathcal{C}_k$ *and* $\mathcal{C}_{k+3 \bmod q}$,

(L5.6.2) $|(\mathcal{C}_k \cup \mathcal{C}_{k+2 \bmod q}) \cap \Omega_{k+1}|$ *and* $|(\mathcal{C}_{k+1 \bmod q} \cup \mathcal{C}_{k+3 \bmod q}) \cap \Omega_{k+2}|$ *are even.*

Then H has a hardness gadget.

Proof. To simplify notation all indices in this proof are considered to be modulo q. We construct a hardness gadget $(I, S, (J_1', y), (J_2', z), (J_3', y, z))$ for H, as defined in Definition 4.6.

Let $\mathcal{C} = (\mathcal{C}_0, \ldots, \mathcal{C}_{q-1})$. Let J_1' and J_2' each be an instance of the generalised cycle gadget $J(J_\mathcal{C}, J_0, \ldots, J_{q-1})$, let $y = v_{k+1}$, and let $z = v_{k+2}$. Then we have $\Omega_y = (\mathcal{C}_k \cup \mathcal{C}_{k+2}) \cap \Omega_{k+1}$ and $\Omega_z = (\mathcal{C}_{k+1} \cup \mathcal{C}_{k+3}) \cap \Omega_{k+2}$ by Lemma 5.5. It follows that $|\Omega_y|$ and $|\Omega_z|$ are even by (L5.6.2). Let $I = \mathcal{C}_{k+2} \cap \Omega_{k+1}$ and $S = \mathcal{C}_{k+1} \cap \Omega_{k+2}$. We note that I and S have odd size by (L5.5.1) and that $I \subset \Omega_y$ and $S \subset \Omega_z$.

Let J_3 be an edge from y to z. For each $o \in \Omega_y \setminus I \subseteq \mathcal{C}_k$, $s \in S \subseteq \mathcal{C}_{k+1}$, $i \in I \subseteq \mathcal{C}_{k+2}$ and $x \in \Omega_z \setminus S \subseteq \mathcal{C}_{k+3}$,

- $|\Sigma_{ox}| = 0$ since no edge exists between $\mathcal{C}_k$ and $\mathcal{C}_{k+3}$ according to (L5.6.1).
- $|\Sigma_{is}| = |\Sigma_{ix}| = |\Sigma_{os}| = 1$ since, by (L5.5.2), for all $\ell \in \{0, \ldots, q-1\}$ we have $\mathcal{C}_\ell \subseteq \Gamma_H(\mathcal{C}_{\ell+1})$. ∎

We point out a corollary which is more easily accessible and does not use the full generality of the gadget $J(J_\mathcal{C}, J_0, \ldots, J_{q-1})$ but rather only uses the cycle gadget $J_\mathcal{C}$.

COROLLARY 5.7. *For an integer $q = 3$ or $q \geq 5$, let H be a graph which contains a cycle $C = c_0, \ldots, c_{q-1}, c_0$ such that*

- *for all $i \in \{0, \ldots, q-1\}$, we have $|N_{C,H}(c_i)| = 1$, and*
- *there is no walk of the form $D = d_0, \ldots, d_{q-1}, d_0$ with $d_i \in \Gamma_H(c_i) \setminus (c_{i-1} \cup c_{i+1})$ ($\forall i \in \{0, \ldots, q-1\}$).*

Then H has a hardness gadget.

Proof. All indices in this proof are considered to be modulo q. For $i \in \{0, \ldots, q-1\}$ we choose $\mathcal{C}_i = N_{C,H}(c_i)$, which by the fact that $|N_{C,H}(c_i)| = 1$ implies $\mathcal{C}_i = \{c_i\}$. We choose $k = 0$. For each $i \in \{0, \ldots, q-1\}$, let (J_i, z_i) be the partially H-labelled graph that only contains the single vertex z_i and has an empty pinning function. It follows that $\Omega_i = V(H)$ and that $J(J_\mathcal{C}, J_0, \ldots, J_{q-1})$ is essentially $J_\mathcal{C}$. We check that the requirements of Lemma 5.6 are met. (L5.5.1) holds since $\mathcal{C}_{i-1} \cap \Omega_i = \mathcal{C}_{i-1} = \{c_{i-1}\}$ and $\mathcal{C}_{i+1} \cap \Omega_i = \mathcal{C}_{i+1} = \{c_{i+1}\}$. (L5.5.2) and (L5.5.3) hold since $|N_{C,H}(c_i)| = 1$ and therefore c_i is the only common neighbour of c_{i-1} and c_{i+1}. There is no walk of the form $D = d_0, \dots, d_{q-1}, d_0$ with $d_i \in \Gamma_H(c_i) \setminus (c_{i-1} \cup c_{i+1})$, as required by (L5.5.4). Since $q \geq 3$ and C is a cycle, the vertices c_0, c_1, c_2 are distinct. If $q = 3$, as C is a cycle, we have $c_0 = c_3$, and (L5.6.1) holds since we do not allow self-loops in H. If otherwise $q \geq 5$ then (L5.6.1) holds since $\Gamma_H(c_1) \cap \Gamma_H(c_3) = N_{C,H}(c_2) = \{c_2\}$ and therefore c_0 (which is a neighbour of c_1) cannot be a neighbour of c_3. Since $q \geq 3$ (L5.6.2) holds as $(\mathcal{C}_0 \cup \mathcal{C}_2) \cap \Omega_1 = \{c_0, c_2\}$ and $(\mathcal{C}_1 \cup \mathcal{C}_3) \cap \Omega_2 = \{c_1, c_3\}$ are sets of 2 distinct vertices.□

Remark 5.8. Suppose that a square-free graph H contains a cycle C. Clearly, the requirements of Corollary 5.7 are met and, by Theorem 4.7, we obtain $\oplus$P-hardness for $\oplus$RET(H). If, in addition, H is involution-free $\oplus$P-hardness carries over to $\oplus$HOM(H) by Theorem 4.4 (from [19, Theorem 3.1]). This argument, together with the classification of $\oplus$HOM(H) for trees by Faben and Jerrum [12] (or alternatively the shorter [19, Lemmas 5.1 and 5.3]) implies the dichotomy for square-free graphs presented in [19].

6. Chordal Bipartite Components. As our main strategy for proving $\oplus$P-hardness of $\oplus$HOM(H) for K_4-minor-free graphs will rely on finding induced cycles whose lengths are not equal to 4. However, this requires us to treat the case of (K_4-minor-free) graphs that include *only* squares as induced cycles separately; recall that such graphs are called chordal bipartite graphs.

In the current section we will construct a hardness gadget for every involution-free, K_4-minor-free, biconnected chordal bipartite graph H, unless H has a very restricted form (this is Lemma 6.17). In this restricted case we call H an *impasse* (which will be formally defined in Definition 6.15). The main tool that we use to construct hardness gadgets relies on two squares that share one edge. More formally, we will consider the following graph:

Definition 6.1 (The graph F, $\Gamma_{H \setminus F}(i, j)$). The graph F is defined to be the graph depicted in Figure 4.

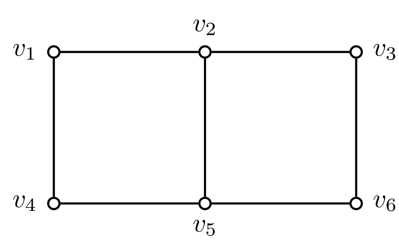

Fig. 4: The graph F.

Given a graph H that contains F as a subgraph, and $i \neq j \in [6]$, we define

$$\Gamma_{H \setminus F}(i, j) = (\Gamma_H(v_i) \cap \Gamma_H(v_j)) \setminus V(F) \,.$$

Definition 6.2 (Type V). Let H be a K_4-minor-free graph that contains F as a subgraph. We say that F has type V in H if one of the following is true

- $\Gamma_{H \setminus F}(1, 5)$ and $\Gamma_{H \setminus F}(3, 5)$ are non-empty and $\Gamma_{H \setminus F}(2, 4)$ and $\Gamma_{H \setminus F}(2, 6)$ are empty.
- $\Gamma_{H \setminus F}(2, 4)$ and $\Gamma_{H \setminus F}(2, 6)$ are non-empty and $\Gamma_{H \setminus F}(1, 5)$ and $\Gamma_{H \setminus F}(3, 5)$ are empty.

An illustration of the former case is given in Figure 5.

The following observation will be useful in the remainder of this section:

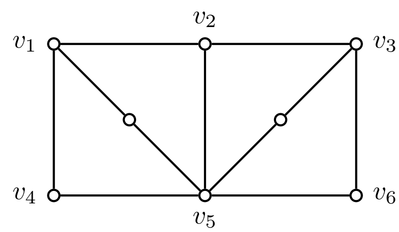

Fig. 5: A K_4-minor-free graph containing F of type V.

LEMMA 6.3. *Let H be a K_4-minor-free graph containing F as a subgraph. At least one of $\Gamma_{H\setminus F}(1,5)$ and $\Gamma_{H\setminus F}(2,4)$ is empty, and at least one of $\Gamma_{H\setminus F}(2,6)$ and $\Gamma_{H\setminus F}(3,5)$ is empty.*

Proof. If $\Gamma_{H\setminus F}(1,5)$ and $\Gamma_{H\setminus F}(2,4)$ are both non-empty, then the vertices v_1, v_2, v_4 and v_5 yield a K_4-minor. If $\Gamma_{H\setminus F}(2,6)$ and $\Gamma_{H\setminus F}(3,5)$ are both non-empty, then the vertices v_2, v_3, v_5 and v_6 yield a K_4-minor. □

LEMMA 6.4. *Let H be a K_4-minor-free graph containing F as a subgraph. If F does not have type V in H then either $\Gamma_{H\setminus F}(1,5) = \Gamma_{H\setminus F}(2,6) = \emptyset$ or $\Gamma_{H\setminus F}(2,4) = \Gamma_{H\setminus F}(3,5) = \emptyset$.*

Proof. Note that either $\Gamma_{H\setminus F}(2,6)$ or $\Gamma_{H\setminus F}(3,5)$ are empty by Lemma 6.3. Assume w.l.o.g. that the former is empty; the other case is symmetric. We distinguish two cases:

(I) $\Gamma_{H\setminus F}(3,5) \neq \emptyset$. Now assume for contradiction that $\Gamma_{H\setminus F}(1,5) \neq \emptyset$. Then, again by Lemma 6.3, we obtain $\Gamma_{H\setminus F}(2,4) = \emptyset$, which implies that F has type V in H, yielding the desired contradiction. In combination with the previous assumption, we thus have $\Gamma_{H\setminus F}(1,5) = \Gamma_{H\setminus F}(2,6) = \emptyset$.

(II) $\Gamma_{H\setminus F}(3,5) = \emptyset$. By Lemma 6.3 we have that either $\Gamma_{H\setminus F}(1,5)$ or $\Gamma_{H\setminus F}(2,4)$ is empty. This concludes the proof as the current case provides additionally $\Gamma_{H\setminus F}(3,5) = \emptyset$ and $\Gamma_{H\setminus F}(2,6) = \emptyset$. □

LEMMA 6.5. *Let H be a K_4-minor-free graph containing F as a subgraph. Then H has a hardness gadget, unless F has type V in H.*

Proof. Using Lemma 6.4 and the fact that H is K_4-minor free, we can w.l.o.g. assume that

(a) The edges $\{v_1, v_6\}$ and $\{v_3, v_4\}$ are *not* present in H as, otherwise, we obtain a K_4-minor.
(b) $\Gamma_H(v_1) \cap \Gamma_H(v_5) = \{v_2, v_4\}$.
(c) $\Gamma_H(v_2) \cap \Gamma_H(v_6) = \{v_3, v_5\}$.

This allows us to construct a hardness gadget:

- $S = \{v_5\}$ and $I = \{v_2\}$.
- J_1 is the graph where y is adjacent to a v_1-pin and a v_5-pin Note that $\Omega_y = \{v_2, v_4\}$ by (b).
- J_2 is the graph where z is adjacent to a v_2-pin and a v_6-pin. Note that $\Omega_z = \{v_3, v_5\}$ by (c).
- J_3 is just the edge $\{y, z\}$.

We have $|\Sigma_{v_4,v_5}| = |\Sigma_{v_5,v_2}| = |\Sigma_{v_2,v_3}| = 1$. Furthermore, $|\Sigma_{v_4,v_3}| = 0$ by (a). □

Definition 6.6 (The graph $S_{k,\ell}$). For positive integers k and ℓ, $S_{k,\ell}$ is the graph depicted in Figure 6.

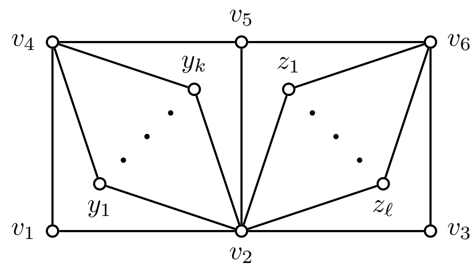

Fig. 6: The graph $S_{k,\ell}$.

LEMMA 6.7. *Let H be a K_4-minor-free graph containing F as a subgraph. If F has type* V *in H and $|\Gamma_{H\setminus F}(1,5)|$ and $|\Gamma_{H\setminus F}(2,4)|$ are even, then H has a hardness gadget.*

Proof. As F has type V in H we can assume w.l.o.g. that $\Gamma_{H\setminus F}(2,4) \neq \emptyset$ and $\Gamma_{H\setminus F}(2,6) \neq \emptyset$, and that $\Gamma_{H\setminus F}(1,5) = \Gamma_{H\setminus F}(3,5) = \emptyset$; the other case is symmetric. In other words, there exist positive integers k and ℓ such that H contains the subgraph $S_{k,\ell}$ (Definition 6.6) with $\Gamma_{H\setminus F}(2,4) = \{y_1, \dots, y_k\}$ and $\Gamma_{H\setminus F}(2,6) = \{z_1, \dots, z_\ell\}$. By the premise of the lemma, k must be even. We will emphasise some crucial properties of H:
(a) $\Gamma_H(v_3) \cap \Gamma_H(v_5) = \{v_2, v_6\}$, since $\Gamma_{H\setminus F}(3,5) = \emptyset$.
(b) v_6 is not adjacent to any vertex in $\{y_1, \dots, y_k, v_1\}$: Assuming otherwise, let $w \in \{y_1, \dots, y_k, v_1\}$ be adjacent to v_6. We obtain the following K_4-minor of H:

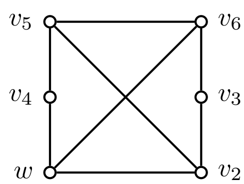

We proceed by constructing a hardness gadget:

- $S = \{v_2\}$ and $I = \{v_5\}$.
- J_1 is the graph where y is adjacent to a v_2-pin and a v_4-pin. Note that

$$\Omega_y = \{v_1, v_5\} \cup \Gamma_{H\setminus F}(2,4) = \{v_1, v_5, y_1, \dots, y_k\}.$$

  In particular, $|\Omega_y|$ is even as k is.
- J_2 is the graph where z is adjacent to a v_3-pin and a v_5-pin. Note that $\Omega_z = \{v_2, v_6\}$ by (a).
- J_3 is just the edge $\{y, z\}$.

We have $|\Sigma_{v_2,v_5}| = |\Sigma_{v_5,v_6}| = 1$ and, for every $o \in \Omega_y \setminus \{v_5\}$, $|\Sigma_{o,v_2}| = 1$. Furthermore, by (b), $|\Sigma_{o,v_6}| = 0$. ∎

6.1. Strong Hardness Gadgets.

Definition 6.8 (strong hardness gadget). A graph J is called a *strong hardness gadget* if every K_4-minor-free graph that contains J as a subgraph has a hardness gadget.

LEMMA 6.9. *The following graph J is a strong hardness gadget:*

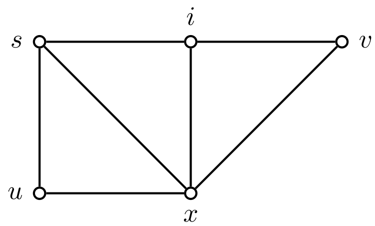

Proof. Let H be a K_4-minor-free supergraph of J. We construct a hardness gadget of H:

- $S = \{s\}$ and $I = \{i\}$.
- J_1 is the graph where y is adjacent to a u-pin and an i-pin. Note that $\Omega_y = \{x, s\}$ as H is K_4-minor free.
- J_2 is the graph where z is adjacent to a v-pin and an s-pin. Note that $\Omega_z = \{x, i\}$ as H is K_4-minor free.
- J_3 is just the edge $\{y, z\}$.

We have $|\Sigma_{s,i}| = |\Sigma_{s,x}| = |\Sigma_{x,i}| = 1$ and $|\Sigma_{x,x}| = 0$ — recall that we do not allow self-loops. ∎

For the proof of the following lemma recall the definition of walk-neighbour-sets from Definition 4.2.

LEMMA 6.10. *Let H be a K_4-minor-free graph containing two adjacent vertices a and b such that $|\Gamma_H(a) \cap \Gamma_H(b)|$ is odd and at least 3. Then H has a hardness gadget.*

Proof. Let c be a common neighbour of a and b and consider the triangle $C = (a, b, c, a)$: If a and c have a common neighbour apart from b, or if b and c have a common neighbour apart from a then Lemma 6.9 applies, as a and b have a common neighbour apart from c by assumption. Otherwise, we have that $|N_{C,H}(a)| = 1$, $|N_{C,H}(b)| = 1$, and $|N_{C,H}(c)| = j \geq 3$, where j is odd. For any $w \in N_{C,H}(c)$ we can assume that

$$\Gamma_H(w) \cap \Gamma_H(a) = \{b\} \text{ and } \Gamma_H(w) \cap \Gamma_H(b) = \{a\}, \tag{6.1}$$

as otherwise we obtain a hardness gadget from Lemma 6.9 (choose w instead of c). Next we can apply Lemma 5.6 to obtain a hardness gadget as follows.

Let $q = 3$ and $\mathcal{C}_0 = N_{C,H}(a) = \{a\}$, $\mathcal{C}_1 = N_{C,H}(b) = \{b\}$, $\mathcal{C}_2 = N_{C,H}(c)$. For each $i \in \{0, 1, 2\}$, let (J_i, z_i) be the partially H-labelled graph that only contains the single vertex z_i and has an empty pinning function. It follows that $\Omega_i = V(H)$. We choose $k = 0$ and check that the requirements of Lemma 5.6 are met.

- (L5.5.1) holds since, for each $i \in \{0, 1, 2\}$, $\Omega_i = V(H)$ and $\mathcal{C}_i$ has odd cardinality (either 1 or j).
- (L5.5.2) and (L5.5.3) hold by (6.1) and the fact that $\Gamma_H(a) \cap \Gamma_H(b) = N_{C,H}(c) = \mathcal{C}_2$.
- Suppose for contradiction that there exists a walk $D = (d_a, d_b, d_c, d_a)$ with $d_a \in \Gamma_H(a) \setminus \{b, c\}$, $d_b \in \Gamma_H(b) \setminus \{a, c\}$ and $d_c \in \Gamma_H(c) \setminus \{a, b\}$. Consequently, as we do not allow self-loops in H, $d_a \neq a$, $d_b \neq b$ and $d_c \neq c$. Then the vertices d_a, a, b, c induce a K_4-minor (where the path from d_a to b goes via d_b, and the path from d_a to c goes via d_c). Hence (L5.5.4) holds.
- Since $\mathcal{C}_0 = \mathcal{C}_{3 \bmod q} = \{a\}$, (L5.6.1) holds by the fact that we do not allow self-loops in H.

- (L5.6.2) holds since $(\mathcal{C}_0 \cup \mathcal{C}_2) \cap \Omega_1 = \mathcal{C}_0 \cup \mathcal{C}_2$, which has cardinality $j+1$ (as we do not allow self-loops in H and therefore $\mathcal{C}_0 = \{a\}$ and $\mathcal{C}_2 = N_{C,H}(c) = \Gamma_H(a) \cap \Gamma_H(b)$ are disjoint), and $j+1$ is even. Analogously, $(\mathcal{C}_1 \cup \mathcal{C}_2) \cap \Omega_1 = \mathcal{C}_1 \cup \mathcal{C}_2$ has even cardinality. ∎

Lemma 6.11. *The following graph J is a strong hardness gadget:*

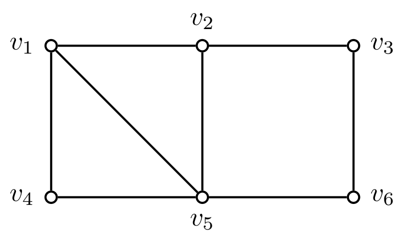

Proof. Let H be a K_4-minor-free supergraph of J. In particular, the graph F is a subgraph of J and thus of H. Note that, due to the edge $\{v_1, v_5\}$, the vertices v_2 and v_4 have no common neighbours apart from v_1 and v_5 in H, as we would obtain a K_4-minor otherwise. In other words, $\Gamma_{H \setminus F}(2,4) = 0$. By Lemma 6.5 we are done, unless F has type V in H. In particular, as $\Gamma_{H \setminus F}(2,4) = \emptyset$, only the following case remains:

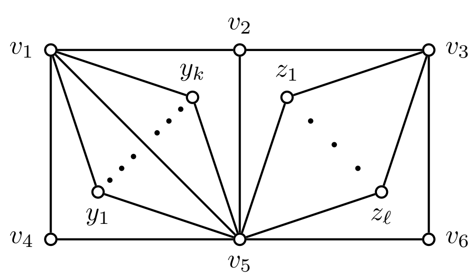

In particular, $\Gamma_{H \setminus F}(1,5) = \{y_1, \ldots, y_k\}$ and $\Gamma_{H \setminus F}(3,5) = \{z_1, \ldots, z_\ell\}$ and $k, \ell > 0$. Now, if k is even, then Lemma 6.7 yields a hardness gadget of H. Finally, if k is odd, then Lemma 6.10 yields a hardness gadget of H — note that Lemma 6.10 is applicable as v_1 and v_5 have precisely $k+2$ common neighbours, which is an odd number greater or equal than 3 since k is odd and positive. ∎

Lemma 6.12. *The following graph J is a strong hardness gadget:*

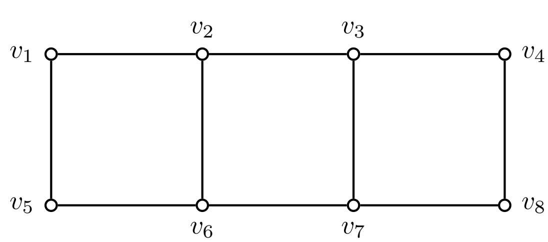

Proof. Let H be a K_4-minor-free supergraph of J.

Claim A *If J is not an induced subgraph of H then H has a K_4-minor or a hardness gadget.*

Proof: If J is not an induced subgraph of H then there is an edge $e = \{v_i, v_j\} \in E(H) \setminus E(J)$ for some $i \neq j \in [8]$. If e is a diagonal of one of the three squares, such as $\{v_2, v_7\}$, then H has a hardness gadget by Lemma 6.11.

If e is not a diagonal of a square, then we obtain a K_4-minor; each case is similar to one of the following two:

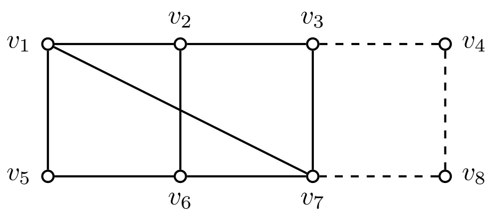

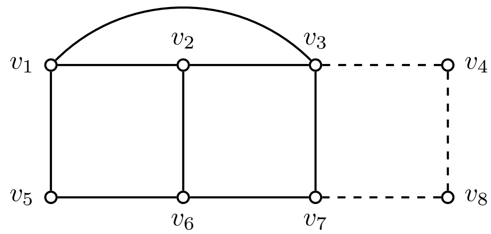

■

Thus assume for the remainder of the proof that J is an induced subgraph of H. Note that J has two subgraphs isomorphic to F. We are done unless both have type V in H by Lemma 6.5. If both have type V, but Lemma 6.7 is applicable, we are done as well. There is thus only one case (up to symmetry) remaining:
(a) $\Gamma_H(v_1) \cap \Gamma_H(v_6) = \{v_2, v_5\}$,
(b) $\Gamma_H(v_3) \cap \Gamma_H(v_6) = \{v_2, v_7\}$,
(c) $\Gamma_H(v_3) \cap \Gamma_H(v_8) = \{v_4, v_7\}$,
(d) $|\Gamma_H(v_2) \cap \Gamma_H(v_5)|$ is odd,
(e) $|\Gamma_H(v_2) \cap \Gamma_H(v_7)|$ is odd, and
(f) $|\Gamma_H(v_4) \cap \Gamma_H(v_7)|$ is odd.
We provide an illustration for convenience:

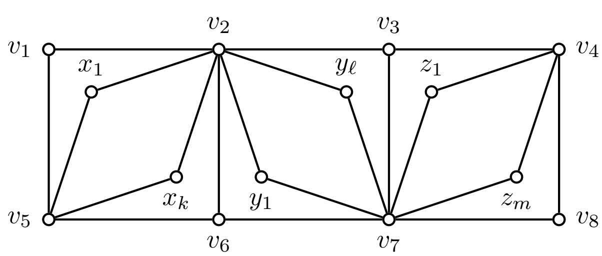

Note that k, ℓ and m are odd. We construct a hardness gadget:

- $S = \{v_2\}$ and $I = \{v_7\}$.
- J_1 is the graph where y is adjacent to a v_1-pin and a v_6-pin. Note that $\Omega_y = \{v_2, v_5\}$ by (a).
- J_2 is the graph where z is adjacent to a v_3-pin and a v_8-pin. Note that $\Omega_z = \{v_7, v_4\}$ by (c).
- J_3 is a path of length 2 from y to z.

By (d), (e) and (f) we have that $|\Sigma_{v_5,v_2}|$, $|\Sigma_{v_2,v_7}|$ and $|\Sigma_{v_7,v_4}|$ are odd. Furthermore, we observe that $|\Sigma_{v_5,v_4}| = 0$ as any path of length 2 from v_5 to v_4 would create a K_4-minor. □

Lemma 6.13. *The following graph J is a strong hardness gadgets:*

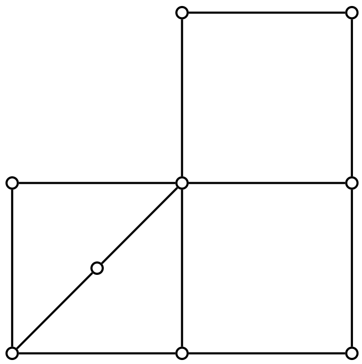

Proof. Let H be a K_4-minor-free supergraph of J. Note that J has two subgraphs isomorphic to F. By Lemma 6.5, we obtain a hardness gadget of H, unless both of the subgraphs isomorphic to F have type V. If this is the case, however, we obtain the following subgraph $\hat{J}$ of H:

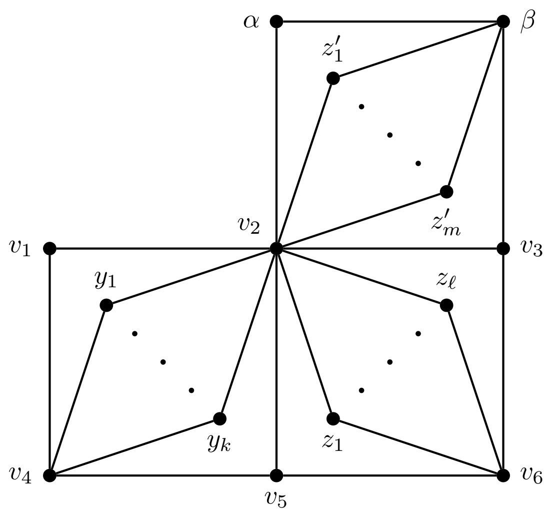

In $\hat{J}$, $k, \ell, m > 0$ and all common neighbours in H between the pairs (v_2, v_4), (v_2, v_6) and (v_2, β) are depicted. By definition of type V, we also obtain that each of the pairs (v_1, v_5), (v_5, v_3) and (v_3, α) has only the two common neighbours in H depicted. Note further, that H has a hardness gadget if at least one of k, ℓ or m is even by Lemma 6.7. Thus assume for the remainder of the proof that all three are odd. We will rely on the following claim, that we can assume $\hat{J}$ to be an *induced* subgraph of H:

Claim A: *If $\hat{J}$ is not an induced subgraph of H then H has a K_4-minor or a hardness gadget.*

Proof: Let $e \in E(H) \setminus E(\hat{J})$ be an edge of H that connects two vertices of $\hat{J}$. We first assume that e connects two vertices in

$$\{v_1, v_2, v_3, v_4, v_5, v_6, y_1, \dots, y_k, z_1, \dots, z_\ell\}\,.$$

We show by case distinction that e either yields a hardness gadget, or a K_4-minor:

(I) $x \in e$ for $x \in \{v_1, y_1, \dots, y_k\}$. Let x' be the other endpoint of e and note that $x' \notin \{v_4, v_2, x\}$ as we do not allow self-loops and multiple edges.

(i) If $x' \in \{v_1, y_1, \dots, y_k, v_5\}$ then we obtain a K_4-minor induced by x, x', v_2, v_4 — note that, as $k > 0$, there exists a 2-path from v_2 to v_4 whose internal vertex is neither x nor x'.

(ii) If $x' \in \{v_3, z_1, \dots, z_\ell\}$ then we obtain a K_4-minor induced by x, v_2, x', v_5 — note that there is a 2-path from v_5 to x via v_4, and a 2-path from v_5 to x' via v_6.

(iii) If $x' = v_6$, then we obtain a K_4-minor induced by x, v_2, v_6, v_5 — note that there is a 2-path from v_5 to x via v_4, and a 2-path from v_6 to v_2 via v_3.

(II) $x \in e$ for $x \in \{v_3, z_1, \dots, z_\ell\}$. Symmetric to the previous case (I).

(III) $v_4 \in e$. Let x' be the other endpoint of e. Note that $x' \notin \{v_4, v_1, y_1, \dots, y_k, v_5\}$ as we do not allow self-loops and multiple edges.

(i) If $x' \in \{v_3, z_1, \dots, z_\ell\}$ then the case is symmetric to case (I)(iii).

(ii) If $x' = v_6$ then we obtain a K_4-minor induced by v_4, v_2, v_6, v_5 — note that there is a 2-path from v_4 to v_2 via v_1, and a 2-path from v_2 to v_6 via v_3.

(iii) If $x' = v_2$, then H has a hardness gadget by Lemma 6.11.

(IV) $v_6 \in e$. Symmetric to the previous case (III).

(V) $v_2 \in e$. Let x' be the other endpoint of e. Then, as we do not allow self-loops and multiple edges, it follows that $x' \notin \{v_2, v_1, y_1, \dots, y_k, v_5, z_1, \dots, z_\ell, v_3\}$. The only remaining candidates for x' are thus v_4 and v_6. However, both of the latter candidates yield a hardness gadget by Lemma 6.11.

(VI) $v_5 \in e$. Let x' be the other endpoint of e and note that $x' \notin \{v_5, v_4, v_2, v_6\}$ as we do not allow self-loops and multiple edges. Similarly as in the previous case (V), all other candidates for x' yield a hardness gadget by Lemma 6.11.

This concludes the case distinction. Observe now, that a symmetric case analysis shows H has a hardness gadget or a K_4-minor if e connects two vertices in

$$\{v_5, v_2, \alpha, v_6, v_3, \beta, z_1, \dots, z_\ell, z'_1, \dots, z'_m\}\,.$$

The remaining possibility for e is to have one endpoint in $\{v_4, v_1, y_1, \dots, y_k\}$ and the other endpoint in $\{\alpha, \beta, z'_1, \dots, z'_m\}$. However, in this case, we find a path from v_5 to v_3 whose vertices are disjoint from $\{v_2, z_1, \dots, z_\ell, v_6\}$. Consequently, we obtain a K_4-minor induced by v_2, v_3, v_5, v_6. ■

We thus assume that $\hat{J}$ is an induced subgraph of H in what follows. Next, we perform a case distinction on the parity of the degree of v_2; in both cases, we construct a hardness gadget.

(I) $\deg_H(v_2)$ is even. We construct a hardness gadget:
- $I = \{v_4\}$ and $S = \{v_6\}$.
- J_1 is the graph where y is adjacent to a v_1-pin and a v_5-pin so $\Omega_y = \{v_2, v_4\}$.
- J_2 is the graph where z is adjacent to a v_5-pin and a v_3-pin so $\Omega_z = \{v_2, v_6\}$.
- J_3 is a 2-path between y and z.

As the degree of v_2 is even, $|\Sigma_{v_2,v_2}|$ is even. As k and ℓ are odd, $|\Sigma_{v_2,v_6}|$ and $|\Sigma_{v_2,v_4}|$ are odd. Finally, we claim that $|\Sigma_{v_6,v_4}|$ is odd: Otherwise there must be an additional 2-path from v_6 to v_4. As $\hat{J}$ is an induced subgraph of H, the internal vertex of this path, let us call it x, cannot be contained in $V(\hat{J})$; otherwise, H would contain an edge between x and a vertex v of $\hat{J}$ while x and v are not adjacent in $\hat{J}$.

This, however, yields a K_4-minor induced by the vertices v_4, v_2, v_6 and v_5 —

note that v_4 and v_2 are connected by the 2-path via v_1, v_2 and v_6 are connected by the 2-path via v_3, and v_4 and v_6 are connected by the 2-path via x.

(II) $\deg_H(v_2)$ is odd. We construct a hardness gadget:
- $I = S = \{v_2\}$.
- J_1 is the graph where y is adjacent to a v_1-pin and a v_5-pin so $\Omega_y = \{v_2, v_4\}$.
- J_2 is the graph where z is adjacent to an α-pin and a v_3-pin so $\Omega_z = \{v_2, \beta\}$.
- J_3 is a 2-path between y and z.

As the degree of v_2 is odd, $|\Sigma_{v_2,v_2}|$ is odd. As k and ℓ are odd, we have that $|\Sigma_{v_4,v_2}|$ and $|\Sigma_{v_2,\beta}|$ are odd. Finally, we claim that $|\Sigma_{v_4,\beta}|$ is even: Assuming otherwise, there must be at least one 2-path in H from v_4 to β; we show that there is none.

As $\hat{J}$ is an induced subgraph of H, the internal vertex of this path, let us call it x, cannot be contained in $V(\hat{J})$; otherwise, H would contain an edge between x and a vertex v of $\hat{J}$ while x and v are not adjacent in $\hat{J}$.

This, however, yields a K_4-minor induced by the vertices v_2, β, v_5 and v_4 — note that v_4 and v_2 are connected by the 2-path via v_1, v_2 and β are connected by the 2-path via α, v_4 and β are connected by the 2-path via x, and v_5 and β are connected by the 3-path via v_6 and v_3. ∎

6.2. Chordal Bipartite Component Lemma.

Definition 6.14 ((1,2)-supergraph). Let J be a connected graph. We say that a supergraph H of J is a *(1,2)-supergraph* of J if every edge of H connecting vertices of J is also an edge of J and every length-2 path of H connecting vertices of J is also a path of J.

For what follows, recall that a chordal bipartite graph is a graph in which every induced cycle is a square. The following notion captures the K_4-minor-free (biconnected) graphs that are obtained by gluing squares together without inducing $\oplus$P-hardness.

Definition 6.15 (impasse, pair of connectors). A K_4-minor-free biconnected graph B is called an *impasse* if there are odd positive integers k and ℓ such that B is a (1,2)-supergraph of the graph $S_{k,\ell}$. Also, with the vertex labels from Definition 6.6, all of the vertices in $\{v_1, y_1, \dots, y_k, v_3, z_1, \dots, z_\ell\}$ are required to have degree 2 in B. The pair (v_1, v_3) is called a *pair of connectors* of the impasse B. (Note that a pair of connectors of B is not unique as, for instance, (v_1, z_1) is also a pair of connectors.)

The graph in Figure 5 is an example of an impasse.

Definition 6.16 (diamond). A biconnected graph B is a *diamond* if, for an integer $k \geq 2$, $V(B) = \{s, t, x_1, \dots, x_k\}$ and $E(B) = \cup_{i \in [k]} \{\{s, x_i\}, \{x_i, t\}\}$.

Note that a square is a diamond with $k = 2$. The following lemma classifies biconnected chordal bipartite graphs:

Lemma 6.17 (Chordal Bipartite Component Lemma). *Let H be a K_4-minor-free graph and let B be a biconnected component of H. If B is chordal bipartite and not just a single edge, then at least one of the following is true:*
(a) B is a diamond.
(b) H has a hardness gadget.

(c) B is an impasse.

Proof. As B is biconnected, chordal bipartite and not a single edge, there exists an induced square $C = (a,b,c,d,a)$ in B. Let us write $\Gamma_{H\setminus C}(a,c)$ for the set $\Gamma_H(a) \cap \Gamma_H(c) \setminus \{b,d\}$ and $\Gamma_{H\setminus C}(b,d)$ for the set $\Gamma_H(b) \cap \Gamma_H(d) \setminus \{a,c\}$. Since B is a biconnected component of H, and $a,b,c,d \in B$, we actually have that $\Gamma_{H\setminus C}(a,c) = \Gamma_B(a) \cap \Gamma_B(c) \setminus \{b,d\}$ and $\Gamma_{H\setminus C}(b,d) = \Gamma_B(b) \cap \Gamma_B(d) \setminus \{a,c\}$. As H is K_4-minor free, we observe that at least one of $\Gamma_{H\setminus C}(a,c)$ and $\Gamma_{H\setminus C}(b,d)$ is empty. Assume w.l.o.g., that $\Gamma_{H\setminus C}(b,d)$ is empty. Let B' be the graph consisting of C together with the edges from a and c to $\Gamma_{H\setminus C}(a,c)$. If $B = B'$ then B is a diamond. Otherwise, as B is biconnected, there is a shortest path P in B connecting two vertices of $C \cup \Gamma_{H\setminus C}(a,c)$ whose internal vertices are not in B'. This path P has an internal vertex since B has no triangle.

Claim A: *P has length 3, one endpoint of P is contained in $\Gamma_B(a) \cap \Gamma_B(c)$ and the other endpoint is contained in $\{a,c\}$.*

Proof: Assume first, for contradiction, that both endpoints of P, let us call them s and t, are in $\Gamma_B(a) \cap \Gamma_B(c)$. The only possible length for P under this assumption is 2, as, otherwise, we obtain an induced cycle (a,s,P,t,a) of length $\neq 4$. As P must have length 2, the endpoints of P cannot be b and d, as $\Gamma_{H\setminus C}(b,d)$ is empty. Thus we can assume w.l.o.g. that $s \neq b$ and $t \neq b$, which yields the following K_4-minor; P is depicted dashed:

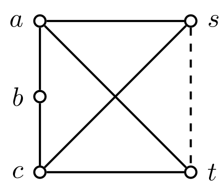

This yields the desired contradiction.

Next, if P starts in a and ends in c, then we obtain an induced cycle that is not a square, unless P as length 2. However, in the latter case, the internal vertex of P is contained in $\Gamma_{H\setminus C}(a,c)$, contradicting the fact that P is not fully contained in $C \cup \Gamma_{H\setminus C}(a,c)$. This shows that one endpoint of P is in $\Gamma_B(a) \cap \Gamma_B(c)$ and the other endpoint is in $\{a,c\}$.

Recall that the length of P is greater than 1 (since it has internal vertices). If P has length 2, then we obtain a triangle, contradicting the fact that B is chordal bipartite. Finally, if P has length at least 4, we obtain an induced cycle of length at least 5, also contradicting chordal-bipartiteness. Consequently, P must have length 3. ■

Claim A yields that B contains a subgraph isomorphic to the graph F — recall from Definition 6.1 that F is just the graph containing two squares that share one edge. We use the vertex labels from Figure 4, i.e., the vertices are $\{v_1,\dots,v_6\}$. Now assume that (b) is not true, i.e., that H does not have a hardness gadget. Using the fact that H is K_4-minor free, and invoking Lemma 6.5 we obtain that F has to be of type V. So, without loss of generality (by renaming), we can assume that $\Gamma_B(v_4) \cap \Gamma_B(v_2) = \{v_1, y_1, \dots, y_k, v_5\}$ and $\Gamma_B(v_6) \cap \Gamma_B(v_2) = \{v_5, z_1, \dots, z_\ell, v_3\}$ for some $k, \ell \geq 1$. Consequently, since H is K_4-minor-free $\Gamma_B(v_1) \cap \Gamma_B(v_5) = \{v_2, v_4\}$ and $\Gamma_B(v_3) \cap \Gamma_B(v_5) = \{v_2, v_6\}$ have to hold. By Lemma 6.7, we obtain that k and ℓ have to be odd. So we have shown that B contains $S_{k,\ell}$ (Definition 6.6) as a subgraph.

Claim B *$S_{k,\ell}$ is an induced subgraph of B.*

Proof: Assume that $S_{k,\ell}$ is not an induced subgraph. Then B (equivalently, H) contains an edge $e \notin E(S_{k,\ell})$ between two vertices of $S_{k,\ell}$. We need to distinguish a variety of (simple) cases:

- $v_4 \in e$: The other endpoint of e cannot be one of $v_4, v_5, y_1, \ldots, y_k, v_1$ as we do not allow self-loops and multi-edges. Further, it cannot be v_6 or v_2, as this would create a triangle, contradicting the fact that B is chordal-bipartite. Finally, if the other endpoint of e is $x \in \{v_3, z_1, \ldots, z_\ell\}$, then we obtain a K_4-minor induced by the vertices v_4, v_5, v_2 and x — note that there is a 2-path from x to v_5 via v_6, and a 2-path from v_2 to v_4 via v_1.
- $v_6 \in e$: Symmetric to the previous case.
- $x \in e$ for some $x \in \{v_1, y_1, \ldots, y_k\}$: The other endpoint of e cannot be one of

  $$v_1, y_1, \ldots, y_k, v_4, v_2, v_5, v_3, z_1, \ldots, z_\ell\,,$$

  as each of those cases would yield a self-loop, a multi-edge, or a triangle (in B). The remaining candidate for the other endpoint is v_6, which was covered in the previous case.
- $x \in e$ for some $x \in \{v_3, z_1, \ldots, z_\ell\}$: Symmetric to the previous case.
- $v_5 \in e$: Any (additional) edge from v_5 to a vertex of $S_{k,\ell}$ would create either a multi-edge, a self-loop, or a triangle.
- $v_2 \in e$: The other endpoint of e cannot be v_2 as we this would create a self-loop. Consequently, one of the previous cases must be true for the other endpoint of e.

■

Recall that we want to show that (a) B is a diamond, (b) H has a hardness gadget, or (c) B is an impasse. For what follows, we distinguish two cases:

(I) All vertices $v_1, y_1, \ldots, y_k, z_1, \ldots, z_\ell, v_3$ have degree 2 in B. In this case we will show that B is a (1,2)-supergraph of $S_{k,\ell}$. This implies (see Definition 6.15) that B is an impasse, so we are finished. To see that B is a (1,2)-supergraph of $S_{k,\ell}$, recall (from Claim B) that $S_{k,\ell}$ is an induced subgraph of B. All neighbours of $v_1, y_1, \ldots, y_k, z_1, \ldots, z_\ell, v_3$ in B are included in $S_{k,\ell}$. Thus, it suffices to show that B has no 2-path connecting vertices in $\{v_4, v_5, v_6, v_2\}$ whose internal vertex x, is outside of $S_{k,\ell}$. We noted above that $\Gamma_B(v_4) \cap \Gamma_B(v_2) \subseteq V(S_{k,\ell})$ and $\Gamma_B(v_6) \cap \Gamma_B(v_2) \subseteq V(S_{k,\ell})$. There is no 2-path in B from v_2 to v_5 because that would yield a triangle in B. Similarly, 2-paths from v_5 to v_4 or v_6 would yield triangles in B, so the only possibility is a 2-path from v_4 to v_6 but this would yield the K_4-minor $\{v_4, v_5, v_6, v_2\}$ in B, contradicting the fact that H (hence B) has no K_4-minor.

(II) Otherwise, assume w.l.o.g. that v_1 has degree at least 3 in B. As B is biconnected, there exists a shortest path P in the remainder of B connecting v_1 with another vertex w of $S_{k,\ell}$. We claim that the only candidates for w are v_4 and v_2, which we will prove by case distinction:

- $w \in \{y_1, \ldots, y_k, v_5\}$. Then we obtain a K_4-minor: (v_4, w, v_2, v_1, v_4) is a square, P connects v_1 and w via vertices not contained in $S_{k,\ell}$, and v_4 and v_2 are connected by a 2-path via a vertex $x \in \{y_1, \ldots, y_k, v_5\} \setminus w$ — note that x exists as $k \geq 1$.
- $w = v_6$. Then we obtain a K_4-minor induced by the vertices v_5, v_6, v_1 and v_2 — note that v_1 is connected to v_5 by the 2-path via v_4, and that v_2 is

connected to v_6 by the 2-path via v_3.

- $w \in \{z_1, \dots, z_\ell, v_3\}$. Then we obtain a K_4-minor induced by the vertices v_5, v_1, v_2 and w — note that v_1 is connected to v_5 by the 2-path via v_4, and that v_5 is connected to w by the 2-path via v_6.

Consequently, w must either be v_4 or v_2 as all other possibilities create a K_4-minor. As B is chordal bipartite, P must have length three. However, if P connects v_1 and v_2, we obtain a strong hardness gadget by Lemma 6.13, and if P connects v_1 and v_4, we obtain a strong hardness gadget by Lemma 6.12. In both cases, H therefore has a hardness gadget. ∎

The following lemma shows that impasses already yield hardness if the vertex v_2 has even degree:

LEMMA 6.18. *Let H be a graph containing an impasse B as biconnected component, that is, there are odd integers k and ℓ such that B is a (1,2)-supergraph of the graph $S_{k,\ell}$ such that, using the vertex labels from Figure 6, all vertices $v_1, y_1, \dots, y_k, v_3, z_1, \dots, z_\ell$ have degree 2 in B. If $\deg_H(v_2)$ is even, then H has a hardness gadget.*

Proof. We construct a hardness gadget:

- $I = \{v_4\}$ and $S = \{v_6\}$.
- J_1 is the graph where y is adjacent to a v_1-pin and a v_5-pin so $\Omega_y = \{v_2, v_4\}$ as H has the impasse B as a biconnected component.
- J_2 is the graph where z is adjacent to a v_5-pin and a v_3-pin so $\Omega_z = \{v_2, v_6\}$ as H has the impasse B as a biconnected component.
- J_3 is a 2-path between y and z.

As the degree of v_2 is even, $|\Sigma_{v_2,v_2}|$ is even. As k and ℓ are odd, we have that $|\Sigma_{v_2,v_4}|$ and $|\Sigma_{v_2,v_6}|$ are odd. Finally, we also have $|\Sigma_{v_4,v_6}| = 1$ as an additional 2-path from v_4 to v_6 would contradict the fact that the biconnected component B of H is an impasse. ∎

7. Sequences of Chordal Bipartite Components.

In Section 6, we proved (Lemma 6.17) that every chordal bipartite biconnected component B of a K_4-minor-free graph H is an edge, a diamond, or an impasse, or the graph H has a hardness gadget. The goal of the current section is to establish a structural property (Lemma 7.14) of a graph H, informally stating that a path in a block-cut-tree of H consisting only of edges, diamonds, and impasses either induces a hardness gadget of H, or has endpoints that satisfy a technical criterion necessary for our construction of global hardness gadgets in Section 9.

Definition 7.1 (good start, good stop). Let H be a graph and let B be a subgraph of H. Let y be a vertex in B and let $L_B \subseteq \Gamma_H(y) \cap V(B)$.

- We say that (L_B, y) is a *good start* in B if there is a gadget (J, z) such that $\{v \in V(H) \mid |\hom((J,z) \to (H,v))| \text{ is odd}\} = L_B \cup R_B$, where $|L_B|$ is odd and $R_B = \Gamma_H(y) \setminus V(B)$.
- We say that (L_B, y) is a *good stop* in B if it is a good start in B and $|R_B|$ is odd.

For non-negative integers k and ℓ, we define some (classes of) graphs with a pair of distinguished vertices a and b each, see Figure 7 (The graph $S_{k,\ell}$ was already defined in Definition 6.15, however, for the scope of this section it will be more convenient to work with the vertex labels as given in Figure 7.).

7.1. Good Starts.

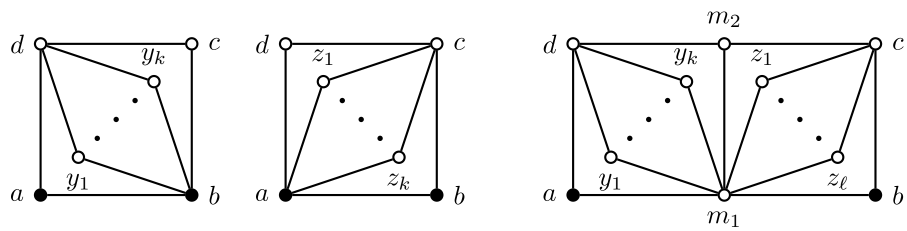

Fig. 7: The graphs BD_k (for "backward diamond"), FD_k (for "forward diamond") and $S_{k,\ell}$ (from left to right).

Lemma 7.2. *Let B be a biconnected component of a graph H, where B is an edge between vertices a and b. Then $(\{a\}, b)$ is a good start in B.*

Proof. Clearly, $\{a\}$ has odd cardinality, and is contained in $\Gamma_H(b) \cap V(B)$. Let (J_B, z_B) be the gadget where z_B is adjacent to a b-pin and let $R_B = \Gamma_H(b) \setminus \{a\}$. Then $\{v \in V(H) \mid |\hom((J_B, z_B) \to (H, v))| \text{ is odd}\} = \Gamma_H(b) = \{a\} \cup R_B$, as desired. ∎

Lemma 7.3. *Let B be a biconnected component of a graph H such that, for an* even *non-negative integer k, B is a graph of the form FD_k and the vertices a and b are as given in Figure 7. Let A be a subgraph of H such that $V(A) \cap V(B) = \{a\}$ Suppose that $L_A \subseteq \Gamma_H(a) \cap V(A)$. If (L_A, a) is a good start in A but not a good stop in A then $(\{a\}, b)$ is a good start in B.*

Proof. By the definition of a good start, $|L_A|$ is odd and there is a gadget (J_A, z_A) such that $\{v \in V(H) \mid |\hom((J_A, z_A) \to (H, v))| \text{ is odd}\} = L_A \cup R_A$ where $R_A = \Gamma_H(a) \setminus V(A)$. Since (L_A, a) is not a good stop in A, $|R_A|$ is even.

Let $L_B = \{a\}$. We now prove the lemma by showing that (L_B, b) is a good start in B. Clearly, $|L_B|$ is odd.

Let (J_B, z_B) be the gadget where z_B is adjacent to the vertex z_A of the gadget J_A and it is also adjacent to a b-pin. In order to prove that (L_B, b) is a good start we check that $\{v \in V(H) \mid |\hom((J_B, z_B) \to (H, v))| \text{ is odd}\} = L_B \cup R_B$, where $L_B = \{a\}$ and $R_B = \Gamma_H(b) \setminus V(B)$.

Since z_B is adjacent to a b-pin we need only consider each $v \in \Gamma_H(b)$ and homomorphisms with $z_B \mapsto v$. Then z_A is also adjacent to z_B and can be mapped to every vertex in the set $\Gamma_H(v) \cap (L_A \cup R_A)$. We determine the cardinality of this set depending on v:

- If $v = a$ then $\Gamma_H(v) \cap (L_A \cup R_A) = L_A \cup R_A$ and $|L_A \cup R_A|$ is odd, as required.
- If $v = c$ (for c as given in Figure 7) then $v = c$ does not have any neighbours in L_A since every path from L_A to B goes through a because B is a biconnected component of H. Hence, $\Gamma_H(v) \cap (L_A \cup R_A) = \Gamma_H(v) \cap R_A$ and $\Gamma_H(v) \cap R_A = \Gamma_H(c) \cap (\Gamma_H(a) \setminus V(A))$ by definition of R_A. Finally, since B is a biconnected component, the vertices of B have no common neighbours outside of B. Thus $\Gamma_H(c) \cap (\Gamma_H(a) \setminus V(A)) = \{d, b, z_1, \ldots, z_k\}$, which has even cardinality, as required (since k is even).
- If $v \in \Gamma_H(b) \setminus V(B)$ then since B is a biconnected component we have $\Gamma_H(v) \cap \Gamma_H(a) = \{b\}$. Consequently, $\Gamma_H(v) \cap (L_A \cup R_A) = \{b\}$, which is odd, as required. ∎

Lemma 7.4. *Let B be a biconnected component of a graph H such that, for a* *non-negative integer k, B is a graph of the form BD_k and the vertices a and b are as given in Figure 7. Let A be a subgraph of H such that $V(A) \cap V(B) = \{a\}$. Suppose that $L_A \subseteq \Gamma_H(a) \cap V(A)$. If (L_A, a) is a good start in A but not a good stop in A then $(\{a\}, b)$ is a good start in B.*

Proof. The proof is analogous to that of Lemma 7.3. We define $L_B = \{a\}$ and use the same gadget and again have to consider each $v \in \Gamma_H(b)$ and homomorphisms with $z_B \mapsto v$ and consequently determine the cardinality of the set $\Gamma_H(v) \cap (L_A \cup R_A)$ depending on v:

- If $v = a$ then $\Gamma_H(v) \cap (L_A \cup R_A) = L_A \cup R_A$ and $|L_A \cup R_A|$ is odd, as required.
- If $v \in \{c, y_1, \ldots y_k\}$ (as given in Figure 7) then v does not have any neighbours in L_A since every path from L_A to B goes through a. Hence, $\Gamma_H(v) \cap (L_A \cup R_A) = \Gamma_H(v) \cap R_A$ and $\Gamma_H(v) \cap R_A = \Gamma_H(v) \cap (\Gamma_H(a) \setminus V(A))$ by definition of R_A. Finally, since B is a biconnected component, the vertices of B have no common neighbours outside of B, $\Gamma_H(v) \cap (\Gamma_H(a) \setminus V(A)) = \{b, d\}$, which has even cardinality, as required.
- If $v \in \Gamma_H(b) \setminus V(B)$ then since B is a biconnected component we have $\Gamma_H(v) \cap \Gamma_H(a) = \{b\}$. Consequently, $\Gamma_H(v) \cap (L_A \cup R_A) = \{b\}$, which is odd, as required. ∎

LEMMA 7.5. *Let B be a biconnected component of a graph H, where B is an impasse (Definition 6.15). Let (a, b) be a pair of connectors of B and let m_1 be the unique common neighbour of a and b in H (see Figure 7). Suppose further that $\deg_H(m_1)$ is odd. Let A be a subgraph of H such that $V(A) \cap V(B) = \{a\}$. Suppose that $L_A \subseteq \Gamma_H(a) \cap V(A)$. If (L_A, a) is a good start in A but not a good stop in A then $(\{m_1\}, b)$ is a good start in B.*

Proof. By the definition of a good start, $|L_A|$ is odd and there is a gadget (J_A, z_A) such that $\{v \in V(H) \mid |\hom((J_A, z_A) \to (H, v))| \text{ is odd}\} = L_A \cup R_A$, where $R_A = \Gamma_H(a) \setminus V(A)$. Since (L_A, a) is not a good stop, $|R_A|$ is even.

Let $L_B = \{m_1\}$. We now prove the lemma by showing that (L_B, b) is a good start in B. Clearly, $|L_B|$ is odd.

Let (J_B, z_B) be the gadget that consists of the gadget J_A joined with a path of length 2 from the vertex z_A to the vertex z_B, and a b-pin that is adjacent to z_B. In order to prove that (L_B, b) is a good start we check that $\{v \in V(H) \mid |\hom((J_B, z_B) \to (H, v))| \text{ is odd}\} = L_B \cup R_B$, where $L_B = \{m_1\}$ and $R_B = \Gamma_H(b) \setminus V(B)$.

Since z_B is adjacent to a b-pin we need only consider $v \in \Gamma_H(b)$ and homomorphisms with $z_B \mapsto v$. Then there is a path of length 2 from z_A to z_B and therefore, for $v \in \Gamma_H(b)$,

$$|\hom((J_B, z_B) \to (H, v))| = |\{u \in L_A \cup R_A \mid \ |\Gamma_H(u) \cap \Gamma_H(v)| \text{ is odd.}\}|.$$

We determine $|\{u \in L_A \cup R_A \mid \ |\Gamma_H(u) \cap \Gamma_H(v)| \text{ is odd.}\}|$ depending on v and using the vertex labels from Figure 7. Note that m_1 and c are the only neighbours of b in B since the degree of b is 2 in B (by the definition of an impasse).

- Consider $v = m_1$.
  - If $u \in \Gamma_H(a) \setminus \{d, m_1\}$ then $u \notin V(B)$ since $\deg_B(a) = 2$. As B is a biconnected component, it follows that a is the only common neighbour of $v = m_1$ and u.
  - The vertices $v = m_1$ and $u = d$ have an odd number of common neighbours in $S_{k,\ell}$ since k is odd. They have no further common neighbours

in B, since B is an impasse, and no further common neighbours in H since B is a biconnected component of H.
  - Finally, $v = m_1$ and $u = m_1$ have an odd number of common neighbours since $\deg_H(m_1)$ is odd by assumption of the lemma.

  Therefore, $\{u \in L_A \cup R_A \mid |\Gamma_H(u) \cap \Gamma_H(v)| \text{ is odd}\} = L_A \cup R_A$ and $|L_A \cup R_A|$ is odd, as required.
- Consider $v = c$.
  - If $u \in \Gamma_H(a) \setminus \{d, m_1\}$, then $u \notin V(B)$ and, as B is a biconnected component, $v = c$ and u have no common neighbours.
  - The vertices $v = c$ and $u = d$ have one common neighbour in $S_{k,\ell}$ (the vertex m_2) and no further common neighbours in H (by the same argument as we used for $v = m_1$), so $v = c$ and $u = d$ have an odd number of common neighbours in H.
  - Finally, $v = c$ and $u = m_1$ have an odd number of common neighbours (since ℓ is odd).

  Therefore, $\{u \in L_A \cup R_A \mid |\Gamma_H(u) \cap \Gamma_H(v)| \text{ is odd}\} = \{d, m_1\}$ which has even cardinality, as required.
- Consider $v \in \Gamma_H(b) \setminus V(B)$.
  - If $u \in \Gamma_H(a) \setminus \{d, m_1\}$ then $u \notin V(B)$ (since $\deg_B(a) = 2$) and, as B is a biconnected component, v and u have no common neighbours.
  - If $u = d$ then $\{u, b\}$ is not an edge of B (by the definition of impasse) so it is not an edge of H (since B is a biconnected component). Hence b is not a common neighbour of u and v. Also, v and u have no other common neighbours since v is not in the biconnected component containing b and d.
  - If $u = m_1$ then the only neighbour of u and v is b since v is not in the biconnected component containing m_1 and b.

  Since $L_A \cup R_A \subseteq \Gamma_H(a)$ and $m_1 \in R_A$ it follows that $\{u \in L_A \cup R_A \mid |\Gamma_H(u) \cap \Gamma_H(v)| \text{ is odd}\} = \{m_1\}$ which has odd cardinality, as required. $\square$

7.2. Good Stops.

Lemma 7.6. *Let B be a biconnected component of a graph H. Suppose that, for an* even *non-negative integer k, B is a graph of the form FD_k with vertices as given in Figure 7. If $(\{a\}, b)$ is a good stop in B then H has a hardness gadget.*

Proof. By the definition of a good stop, $R_B = \Gamma_H(b) \setminus \{a, c\}$ has odd cardinality and there is a gadget (J_B, z_B) such that $\{v \in V(H) \mid |\hom((J_B, z_B) \to (H, v))| \text{ is odd}\}$ equals $\{a\} \cup R_B$.

We give a hardness gadget $(I, S, (J_1, y), (J_2, z), (J_3, y, z))$ for H as follows:

- $I = \{a\}$ and $S = \{b\}$.
- J_1 is the gadget J_B with $y = z_B$ so $\Omega_y = \{a\} \cup R_B$, which has even cardinality, as required.
- J_2 is the graph where z is adjacent to an a-pin and a c-pin so we have $\Omega_z = \{b, d, z_1, \ldots, z_k\}$, which has even cardinality, as required (since k is even).
- J_3 is an edge between y and z.

Note that a is adjacent to every vertex in Ω_z, and b is adjacent to every vertex in Ω_y, as required. Since $\Omega_y \setminus I = R_B = \Gamma_H(b) \setminus \{a, c\}$ and $\Omega_z \setminus S = \{d, z_1, \ldots, z_k\}$ and B is a biconnected component, there is no edge from $\Omega_y \setminus I$ to $\Omega_z \setminus S$, as required. $\square$

Lemma 7.7. *Let B be a biconnected component of a graph H. Suppose that, for* *a non-negative integer k, B is a graph of the form BD_k with vertices as given in Figure 7. If $(\{a\}, b)$ is a good stop in B then H has a hardness gadget.*

Proof. By the definition of a good stop, $R_B = \Gamma_H(b) \setminus V(B)$ has odd cardinality and there is a gadget (J_B, z_B) such that $\{v \in V(H) \mid |\hom((J_B, z_B) \to (H, v))| \text{ is odd}\}$ equals $\{a\} \cup R_B$. We give a hardness gadget $(I, S, (J_1, y), (J_2, z), (J_3, y, z))$ for H as follows:

- $I = \{a\}$ and $S = \{b\}$.
- J_1 is the gadget J_B with $y = z_B$ so $\Omega_y = \{a\} \cup R_B$, which has even cardinality, as required.
- J_2 is the graph where z is adjacent to an a-pin and a c-pin so $\Omega_z = \{b, d\}$, which has even cardinality, as required.
- J_3 is an edge between y and z.

Note that a is adjacent to every vertex in Ω_z, and b is adjacent to every vertex in Ω_y, as required. Since $R_B = \Gamma_H(b) \setminus V(B)$ and B is a biconnected component, note that there are no edges between $\Omega_y \setminus I = R_B$ and $\Omega_z \setminus S = \{d\}$, as required. ∎

LEMMA 7.8. *Let B be a biconnected component of a graph H. Suppose that B is an impasse (Definition 6.15) and that (a, b) is a pair of connectors of B. Let m_1 be the unique common neighbour of a and b in H (see Figure 7). Suppose further that $\deg_H(m_1)$ is odd. If $(\{m_1\}, b)$ is a good stop in B then H has a hardness gadget.*

Proof. By the definition of a good stop, $R_B = \Gamma_H(b) \setminus V(B)$ has odd cardinality and there is a gadget (J_B, z_B) such that $\{v \in V(H) \mid |\hom((J_B, z_B) \to (H, v))| \text{ is odd}\}$ equals $\{m_1\} \cup R_B$, Using the vertex labels from Figure 7, we give a hardness gadget $(I, S, (J_1, y), (J_2, z), (J_3, y, z))$ for H as follows:

- $I = \{m_1\}$ and $S = \{m_1\}$.
- J_1 is the gadget J_B with $y = z_B$ so $\Omega_y = \{m_1\} \cup R_B$, which has even cardinality, as required.
- J_2 is the graph where z is adjacent to an a-pin and an m_2-pin so $\Omega_z = \{m_1, d\}$, which has even cardinality, as required.
- J_3 is a 2-path between y and z.

There are an odd number of 2-walks from m_1 to itself since $\deg_H(m_1)$ is odd by assumption. There are an odd number of 2-walks from m_1 to d since k is odd and no pair of vertices of $S_{k,\ell}$ has common neighbours outside of $S_{k,\ell}$. Since $R_B = \Gamma_H(b) \setminus V(B)$ and B is biconnected there is exactly one 2-walk from m_1 to each vertex in R_B. Thus, for $s \in S = \{m_1\}$, $i \in I = \{m_1\}$, $o \in \Omega_y \setminus I = R_B$, $x \in \Omega_z \setminus S = \{d\}$, we have shown that $|\Sigma_{i,s}|$, $|\Sigma_{o,s}|$ and $|\Sigma_{i,x}|$ are odd, as required. Finally, since B is a biconnected component there are no 2-walks from d to a vertex in R_B and therefore $|\Sigma_{o,x}|$ is even, as required. ∎

7.3. Hardness Results. In this section we establish hardness results which are used to prove Lemma 7.14 in Section 7.4.

LEMMA 7.9. *Let H be a graph and let B be a biconnected component of H. Suppose that, for an* odd *non-negative integer k, B is a graph of the form FD_k with vertex labels as given in Figure 7. If $\deg_H(a)$ is even then H has a hardness gadget.*

Proof. We give a hardness gadget $(I, S, (J_1, y), (J_2, z), (J_3, y, z))$ for H as follows:

- $I = \{a\}$ and $S = \{b, d, z_1, \ldots z_k\}$.
- J_1 is the graph where y is adjacent to a b-pin and a d-pin so $\Omega_y = \{a, c\}$, which has even cardinality, as required.
- J_2 is the graph where z is adjacent to an a-pin so $\Omega_z = \Gamma_H(a)$, which has

even cardinality, as required.

- J_3 is an edge between y and z.

Note that a is adjacent to every vertex in Ω_z, and each vertex of S is adjacent to every vertex in Ω_y, as required. Since B is a biconnected component there are no edges between $\Omega_y \setminus I = \{c\}$ and $\Omega_z \setminus S = \Gamma_H(a) \setminus V(B)$, as required. ∎

LEMMA 7.10. *Let H be a graph and let A and B be biconnected components of H. Suppose that, for* odd *integers $k \geq 1$ and $\ell \geq 1$, there is an isomorphism f from the graph BD_k to A and an isomorphism g from the graph FD_ℓ to B. Suppose that there is a vertex $w = f(b) = g(a)$ such that $\deg_H(w)$ is odd. Then H has a hardness gadget.*

Proof. We give a hardness gadget $(I, S, (J_1, y), (J_2, z), (J_3, y, z))$ for H as follows:

- $I = \{w\}$ and $S = \{w\}$.
- J_1 is the graph where y is adjacent to an $f(a)$-pin and an $f(c)$-pin so $\Omega_y = \{f(d), f(b)\} = \{f(d), w\}$, which has even cardinality, as required.
- J_2 is the graph where z is adjacent to a $g(b)$-pin and an $g(d)$-pin so $\Omega_z = \{g(c), g(a)\} = \{g(c), w\}$, which has even cardinality, as required.
- J_3 is a 2-path between y and z.

By the fact that A and B are biconnected components, there are exactly $k + 2$ walks of length 2 from $f(d)$ to w, and there are exactly $\ell + 2$ walks of length 2 from $g(c)$ to w, where k and ℓ are odd. Since $\deg_H(w)$ is odd, there is a an odd number of length-2 walks from w to itself. Finally, there are no length-2 walks from $f(d)$ to $g(c)$, as required. ∎

LEMMA 7.11. *Let H be a graph and let B be a biconnected component of H that is of the form BD_k for some integer $k \geq 0$. Using the vertex names from Figure 7, there is a gadget (J, z) such that $\{v \in V(H) \mid |\hom((J, z) \to (H, v))| \text{ is odd}\} = \Gamma_H(b) \setminus V(B)$.*

Proof. The graph J has three pinned vertices — an a-pin, a b-pin, and a c-pin. The b-pin is adjacent to the vertex z and the other two pins are attached to z by paths of length 2.

We will now consider each $v \in V(H)$ to determine whether $|\hom((J, z) \to (H, v))|$ is odd. Since z is adjacent to a b-pin in J, this can only be true for $v \in \Gamma_H(b)$.

First, consider a vertex $v \in \Gamma_H(b) \cap V(B)$.

- If $v \in \{a, y_1, \ldots, y_k\}$ then v has exactly two length-2 walks to c, therefore $|\hom((J, z) \to (H, v))|$ is even.
- If $v = c$ then v has exactly two length-2 walks to a so $|\hom((J, z) \to (H, v))|$ is even.

Finally, consider a vertex $v \in \Gamma_H(b) \setminus V(B)$. There is exactly one 2-walk to a, and exactly one 2-walk to c, so $|\hom((J, z) \to (H, v))|$ is odd. ∎

The following lemma is essentially the same as Lemma 7.11.

LEMMA 7.12. *Let H be a graph and let B be a biconnected component of H that is of the form FD_k for some integer $k \geq 0$. Using the vertex names from Figure 7, there is a gadget (J, z) such that $\{v \in V(H) \mid |\hom((J, z) \to (H, v))| \text{ is odd}\} = \Gamma_H(a) \setminus V(B)$.*

Proof. The graph J has three pinned vertices — an a-pin, a b-pin, and a d-pin. The a-pin is adjacent to the vertex z and the other two pins are attached to z by paths of length 2.

We will now consider each $v \in V(H)$ to determine whether $|\hom((J, z) \to (H, v))|$

is odd. Since z is adjacent to an a-pin in J, this can only be true for $v \in \Gamma_H(a)$.

First, consider a vertex $v \in \Gamma_H(a) \cap V(B)$.

- If $v \in \{d, z_1, \dots, z_k\}$ then v has exactly two length-2 walks to b, therefore $|\hom((J,z) \to (H,v))|$ is even.
- If $v = b$ then v has exactly two length-2 walks to d so $|\hom((J,z) \to (H,v))|$ is even.

Finally, consider a vertex $v \in \Gamma_H(b) \setminus V(B)$. There is exactly one 2-walk to b, and exactly one 2-walk to d, so $|\hom((J,z) \to (H,v))|$ is odd. ∎

We obtain the following lemma, which is a generalisation of [19, Lemma 4.5].

LEMMA 7.13. *For an integer $q \geq 1$, let $P = v_0, \dots, v_q$ be a path in a graph H. Suppose that no edge of P is part of a square in H and that $\deg_H(v_j)$ is odd for all $j \in [q-1]$. Suppose that*

1 a) $\deg_H(v_0)$ *is even, or*

1 b) $\deg_H(v_0)$ *is odd and there is a biconnected component B_0 that is isomorphic to BD_k for some* odd *integer $k \geq 1$, where the isomorphism maps v_0 to the vertex b from Figure 7.*

Suppose further that

2 a) $\deg_H(v_q)$ *is even, or*

2 b) $\deg_H(v_q)$ *is odd and there is a biconnected component B_{q+1} that is isomorphic to FD_k for some* odd *integer $k \geq 1$, where the isomorphism maps v_q to the vertex a from Figure 7.*

Then H has a hardness gadget.

Proof. We give a hardness gadget $(I, S, (J_1, y), (J_2, z), (J_3, y, z))$ for H as follows:

- $I = \{v_1\}$ and $S = \{v_{q-1}\}$.
- If 1 a) holds, then J_1 is the graph where y is adjacent to a v_0-pin so $\Omega_y = \Gamma_H(v_0)$, which has even cardinality as required. If 1 b) holds then (J_1, y) is the gadget from Lemma 7.11 and $\Omega_y = \Gamma_H(v_0) \setminus V(B_0)$, which has even cardinality as required. The vertex v_1 is in Ω_y because the edge $\{v_0, v_1\}$ is not part of a square in H.
- If 2 a) holds, then J_2 is the graph where z is adjacent to a v_q-pin so $\Omega_z = \Gamma_H(v_q)$, which has even cardinality as required. If 2 b) holds then (J_2, z) is the gadget from Lemma 7.12 and $\Omega_z = \Gamma_H(v_q) \setminus V(B_{q+1})$, which has even cardinality as required. The vertex v_{q-1} is in Ω_z because the edge $\{v_{q-1}, v_q\}$ is not part of a square in H.
- J_3 is the path gadget J_P.

This is a hardness gadget by Lemma 5.2. ∎

7.4. Chordal Bipartite Sequence Lemma.

LEMMA 7.14 (Chordal Bipartite Sequence Lemma). *For an integer $q \geq 1$, let $B_1, \dots, B_q$ be biconnected components of a graph H and let $b_0, \dots, b_q$ be vertices such that, for all $i \in [q]$, b_{i-1} and b_i are distinct vertices of B_i, and B_i satisfies one of the following:*

- *B_i is an edge from b_{i-1} to b_i,*
- *B_i is a diamond in which $\{b_{i-1}, b_i\}$ is an edge, or*
- *B_i is an impasse, where (b_{i-1}, b_i) is a pair of connectors of B_i. In this case, let d_i be the unique common neighbour of b_{i-1} and b_i in H.*

If $|\Gamma_H(b_0) \setminus V(B_1)|$ is odd, then at least one of the following holds:

- *B_q is an edge or a diamond and $(\{b_{q-1}\}, b_q)$ is a good start in B_q but not a good stop in B_q,*

- *B_q is an impasse and $(\{d_q\}, b_q)$ is a good start in B_q but not a good stop in B_q, or*
- *H has a hardness gadget.*

Proof. We start by collecting some facts that we will need.

Fact 1. *If $i \in [q]$ and B_i is a diamond, then at least one of the following holds:*

- *for some non-negative integer k there is an isomorphism from FD_k to B_i, mapping the vertex a from Figure 7 to b_{i-1} and the vertex b to b_i (we refer to this situation below by saying "B_i is of the form FD_k"), or*
- *for some non-negative integer k there is an isomorphism from BD_k to B_i, mapping the vertex a from Figure 7 to b_{i-1} and vertex b to b_i. (We refer to this situation as "B_i is of the form BD_k").*

Fact 2. *If B_1 is an edge or a biconnected component of the form FD_k for an odd integer k then $\Gamma_H(b_0)$ is even.* (This is because b_0 has and odd number of neighbours in B_1 and an odd number outside of B_1, by assumption.)

Let $L_0 = \Gamma_H(b_0) \setminus V(B_1)$ and let B_0 be the subgraph of H induced by the vertices in $L_0 \cup \{b_0\}$. For every $i \in [q]$ such that B_i is an edge or a diamond, let $L_i = \{b_{i-1}\}$. For every $i \in [q]$ such that B_i is an impasse, let $L_i = \{d_i\}$. For every $i \in \{0, \ldots, q\}$, let $R_i = \Gamma_H(b_i) \setminus V(B_i)$.

We start by considering $i \in \{0, 1\}$ and showing that (L_i, b_i) is a good start in B_i or that H has a hardness gadget. We first deal with the easy case $i = 0$. $|L_0|$ is odd by assumption so, to show that (L_0, b_0) is a good start in B_0, it suffices to use the gadget (J, z) in which z is adjacent to a b_0-pin. Note that $R_0 = \Gamma_H(b_0) \cap V(B_1)$. We next deal with $i = 1$ by considering four cases depending on B_1:

- If B_1 is an edge from b_0 to b_1 then, by Lemma 7.2, (L_1, b_1) is a good start in B_1.
- If B_1 is a diamond of the form FD_k for an odd integer $k \geq 0$ then H has a hardness gadget by Fact 2 and Lemma 7.9.
- If B_1 is a diamond of the form FD_k for an even integer $k \geq 0$ then R_0 has even cardinality. Therefore (L_0, b_0) is not a good stop in B_0 and we can apply Lemma 7.3 to show that (L_1, b_1) is a good start in B_1.
- If B_1 is a diamond of the form BD_k then R_0 has even cardinality. Therefore (L_0, b_0) is not a good stop in B_0 and we can apply Lemma 7.4 to show that (L_1, b_1) is a good start in B_1.
- If B_1 is an impasse where (b_0, b_1) is a pair of connectors. Then b_0 has 2 neighbours in B_1 and hence R_0 has even cardinality. Therefore (L_0, b_0) is not a good stop in B_0. Recall that d_1 be the unique common neighbour of b_0 and b_1 in H. If $\deg_H(d_1)$ is even then H has a hardness gadget by Lemma 6.18. Otherwise Lemma 7.5 shows that (L_1, b_1) is a good start in B_1.

For the rest of the proof, let j be the smallest index in $[q]$ that satisfies one of the following properties:

(P1) $|R_j|$ is odd and there is no odd integer k such that B_j is of the form FD_k.
(P2) There is an odd integer k such that B_j is of the form FD_k.
(P3) $j = q$ and $|R_j|$ is even and there is no odd integer k such that B_j is of the form FD_k.

We will use the following claims.

Claim A *Suppose that j does not satisfy (P2). Then H has a hardness gadget or, for all $\ell \in [j]$, the following are satisfied.*

- *(L_ℓ, b_ℓ) is a good start in B_ℓ, and*
- *If $\ell > 1$ then $(L_{\ell-1}, b_{\ell-1})$ is not a good stop in $B_{\ell-1}$.*

Proof: The proof of Claim A is by induction on ℓ. We have already established the base case $\ell = 1$. Now fix $\ell \in \{2, \dots, j\}$ and suppose (from the inductive hypothesis) that $(L_{\ell-1}, b_{\ell-1})$ is a good start in $B_{\ell-1}$. By the minimality of j, $B_{\ell-1}$ is not of the form FD_k for an odd integer k (otherwise $\ell - 1$ would satisfy (P2)). Again, by minimality of j, $|R_{\ell-1}|$ is even (otherwise $\ell-1$ would satisfy (P1)). By the definition of good stop, $(L_{\ell-1}, b_{\ell-1})$ is not a good stop in $B_{\ell-1}$. Since j does not satisfy (P2), B_ℓ is not of the form FD_k for an odd integer k. Thus, we can apply one of Lemmas 7.2, 7.3, 7.4 or 7.5 depending on the form of B_ℓ to show that (L_ℓ, b_ℓ) is a good start in B_ℓ. This completes the proof of Claim A. ■

Claim B *Suppose that B_j satisfies one of the following.*
(B1) B_j is the edge $\{b_{j-1}, b_j\}$ and $\deg_H(b_j)$ is even, or
(B2) there is an odd integer k such that B_j is of the form FD_k. □
Suppose that there is an integer ℓ in the range $1 \leq \ell \leq j$ such that, for $i \in \{\ell, \dots, j-1\}$, B_i is the edge $\{b_{i-1}, b_i\}$ and for $i \in \{\ell, \dots, j\}$, $\Gamma_H(b_{i-1})$ and $\Gamma_H(b_{i-1}) \cap V(B_{i-1})$ have odd cardinality. Then H has a hardness gadget or there is an integer p in the range $1 \leq p \leq \ell - 1$ and an odd integer k' such that B_p is of the form $BD_{k'}$. Also, for $i \in \{p+1, \dots, j-1\}$, B_i is the edge $\{b_{i-1}, b_i\}$ where $\Gamma_H(b_{i-1})$ has odd cardinality.

Proof: The proof of Claim B is by induction on ℓ. The base case $\ell = 1$ is vacuous — taking $i = 1$, the precondition of the claim ensures that $|\Gamma_H(b_0)|$ is odd, contrary to Fact 2. So consider some $\ell > 1$ for which we wish to prove the claim. Since taking $i = \ell$ guarantees that $|\Gamma_H(b_{\ell-1}) \cap V(B_{\ell-1})|$ is odd, $B_{\ell-1}$ is either an edge or it is of the form $BD_{k'}$ for an odd integer k'. We consider each case.

- $B_{\ell-1}$ is an edge: If $\deg_H(b_{\ell-2})$ is even H has a hardness gadget by Lemma 7.13 (take $v_0, \dots, v_q = b_{\ell-2}, \dots, b_j$ in Case (B1) and $v_0, \dots, v_q = b_{\ell-2}, \dots, b_{j-1}$ in Case (B2). Thus, assume that $\deg_H(b_{\ell-2})$ is odd. By Fact 2, $\ell - 2 \geq 1$ and consequently (by Claim A), $(L_{\ell-2}, b_{\ell-2})$ is a good start in $B_{\ell-2}$ that is not a good stop in $B_{\ell-2}$. This implies that $|R_{\ell-2}|$ is even, which together with the fact that $\deg_H(b_{\ell-2})$ is odd implies that $b_{\ell-2}$ has an odd number of neighbours in $B_{\ell-2}$. So the preconditions of the claim are met with $i = \ell - 1$ and we can finish by induction.
- $B_{\ell-1}$ is of the form $BD_{k'}$ for an odd integer k'. The claim follows by taking $p = \ell - 1$.

This concludes the proof of Claim B. ■

We now make a case distinction, depending on which property j satisfies.

Case (P1). We will show that H has a hardness gadget.

By Claim A, either H has a hardness gadget (in which case we are finished) or (L_j, b_j) is a good start in B_j. Since $|R_j|$ is odd, (L_j, b_j) is a good stop in B_j. We now distinguish several cases, depending on the form of B_j.

- If B_j is of the form FD_k for even k, or of the form BD_k, then H has a hardness gadget by Lemmas 7.6 or 7.7, respectively.
- If B_j is an impasse, then depending on the degree of d_j, H has a hardness gadget either by Lemma 6.18 (if the degree of d_j is even) or by Lemma 7.8 (if the degree of d_j is odd).
- Finally, suppose that B_j is an edge. We will use Claim B with $\ell = j$ to show that H has a hardness gadget. The first step is to show that (unless H has a hardness gadget) the preconditions of the claim are met — that is $\deg_H(b_j)$ is even, $\deg_H(b_{j-1})$ is odd, and $|\Gamma_H(b_{j-1}) \cap V(B_{j-1})|$

is odd.

By (P1), $|R_j|$ is odd. Since b_j has only one neighbour in B_j, $\deg_H(b_j)$ is even. If $\deg_H(b_{j-1})$ is even, H has a hardness gadget by Lemma 7.13 (taking $q = 1$, $v_0 = b_{j-1}$ and $v_1 = b_j$). From now on, we assume that $\deg_H(b_{j-1})$ is odd. By Fact 2, $j - 1 \geq 1$. By the minimality of j, $|R_{j-1}|$ is even, which implies that b_{j-1} has an odd number of neighbours in B_{j-1}.

Applying Claim B with $\ell = j$, either H has a hardness gadget. or there is an integer p in the range $1 \leq p \leq j-1$ and an odd integer k' such that B_p is of the form $BD_{k'}$. Also, for $i \in \{p+1, \ldots, j-1\}$, B_i is the edge $\{b_{i-1}, b_i\}$ where $\Gamma_H(b_{i-1})$ has odd cardinality.

Now we apply Lemma 7.13 with the path $v_0, \ldots, v_q$ equal to $b_p, \ldots, b_j$. The degrees of $v_0, \ldots, v_{q-1}$ are odd and the degree of v_q is even. v_0 is in the biconnected component B_p. This shows that H has a hardness gadget.

Case (P2). We will use Claim B with $\ell = j$ to show that H has a hardness gadget. The first step is to show that (unless H has a hardness gadget) the preconditions of the claim are met — that is $\deg_H(b_{j-1})$ is odd, and $|\Gamma_H(b_{j-1}) \cap V(B_{j-1})|$ is odd.

If $\deg_H(b_{j-1})$ is even then H has a hardness gadget by Lemma 7.9. From now on, we assume that $\deg_H(b_{j-1})$ is odd. By Fact 2, $j - 1 \geq 1$. By the minimality of j, $|R_{j-1}|$ is even, which implies that b_{j-1} has an odd number of neighbours in B_{j-1}.

Applying Claim B with $\ell = j$, either H has a hardness gadget. or there is an integer p in the range $1 \leq p \leq j-1$ and an odd integer k' such that B_p is of the form $BD_{k'}$. Also, for $i \in \{p+1, \ldots, j-1\}$, B_i is the edge $\{b_{i-1}, b_i\}$ where $\Gamma_H(b_{i-1})$ has odd cardinality.

If $p = j-1$ then H has a hardness gadget by Lemma 7.10. Otherwise, we apply Lemma 7.13 with the path $v_0, \ldots, v_q$ equal to $b_p, \ldots, b_{j-1}$. The degrees of $v_0, \ldots, v_q$ are odd. v_0 is in the biconnected component B_p and b_q is in the biconnected component B_j. This shows that H has a hardness gadget.

Case (P3) By Claim A, H has a hardness gadget or (L_q, b_q) is a good start in B_q. In the latter case, since $|R_q|$ is even, (L_q, b_q) is not a good stop in B_q.

8. K_4-minor-free Components. In this section, we establish a structural classification for biconnected K_4-minor-free graphs. Recall that such a classification has already been achieved for biconnected *chordal bipartite* graphs in Section 6 (see Lemma 6.17). For this reason, we focus in what follows on biconnected K_4-minor-free graphs that are not chordal bipartite, which is equivalent to focusing on biconnected K_4-minor-free graphs which have an induced cycle that is not a square. The main result of this section is presented in Lemma 8.10. Informally, it states that every biconnected K_4-minor-free graph B either induces a hardness gadget for every K_4-minor-free graph containing B as a biconnected component, or B is an edge, a diamond, an impasse, or a so-called obstruction (Definition 8.6). Together with our insights from the previous section, obstructions will be the final building block in our construction of global hardness gadgets in Section 9.

Definition 8.1 (separation, separator). Let G be a graph and let A and B be subsets of $V(G)$. The pair (A, B) is a *separation* of G if $V(G) = A \cup B$ and G has no edges between $A \setminus B$ and $B \setminus A$. The set $A \cap B$ is called the *separator* of this separation.

8.1. Induced Cycles. Recall Definition 4.2, which defines for a closed walk $W = (w_0, \dots, w_{q-1}, w_0)$ in a graph H the walk-neighbour-set $N_{W,H}(w_i) = \Gamma_H(w_{i-1}) \cap \Gamma_H(w_{i+1})$, where the indices are taken modulo q. In this section we will use this notion mainly for cycles.

LEMMA 8.2. *Let H be a biconnected K_4-minor-free graph containing an induced cycle $C = (c_0, \dots, c_{q-1}, c_0)$ for some $q \neq 4$. Then the walk-neighbour-sets $N_{C,H}(c_0)$, $\dots, N_{C,H}(c_{q-1})$ are pairwise disjoint.*

Proof. If $q = 3$ then the fact that we do not allow self-loops in H together with the fact that H does not contain K_4 as a subgraph ensures that the $N_{C,H}(c_i)$ are pairwise disjoint.

Suppose $q > 4$. Assume for contradiction that there exists a vertex $w \in N_{C,H}(c_i) \cap N_{C,H}(c_j)$ for some $i \neq j$. If w is part of the cycle C, then we obtain a chord (note that $q > 4$), contradicting the fact that C is induced. If w is not part of the cycle C, then w is adjacent to at least 3 vertices of the cycle, yielding a K_4-minor. ∎

LEMMA 8.3. *Let H be a biconnected K_4-minor-free graph containing an induced cycle $C = (c_0, \dots, c_{q-1}, c_0)$. If $q > 4$ and $|N_{C,H}(c_i)| > 1$ for some $i \in \{0, \dots, q-1\}$ then there exists a separation (A, B) of H such that $C \setminus \{c_i\} \subseteq A$, $N_{C,H}(c_i) \subseteq B$ and $A \cap B = \{c_{i-1}, c_{i+1}\}$. Furthermore, H is a (1,2)-supergraph of $H[B]$.*

Proof. Let $S_1, \dots, S_k$ be the connected components of the graph obtained from H by deleting c_{i-1}, c_{i+1}, and all edges incident to c_{i-1} and c_{i+1}. Then w.l.o.g. we can assume that $c_j \in V(S_1)$ for all $j \notin \{i-1, i, i+1\}$. Set $A = V(S_1) \cup \{c_{i-1}, c_{i+1}\}$ and $B = V(S_2) \cup \dots \cup V(S_k) \cup \{c_{i-1}, c_{i+1}\}$. By construction, we have $A \cup B = V(H)$ and $A \cap B = \{c_{i-1}, c_{i+1}\}$, and there are no edges between $A \setminus B$ and $B \setminus A$. We claim that $N_{C,H}(c_i) \cap V(S_1) = \emptyset$ as, otherwise, we obtain the following K_4-minor; recall that $|N_{C,H}(c_i)| > 1$ and $q > 4$:

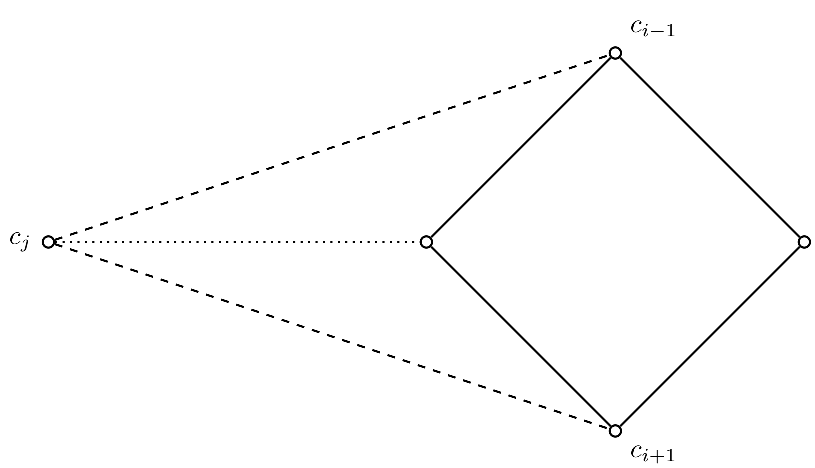

Here, the dashed lines depict the path $c_{i-1}, \dots, c_j, \dots, c_{i+1}$ which is $C \setminus \{c_i\}$. Further, c_j is a vertex of C satisfying that there exists a (shortest) path P from a vertex in $N_{C,H}(c_i)$ to c_j such that the internal vertices of P are disjoint from $C \cup N_{C,H}(c_i)$. Note that c_j exists if $N_{C,H}(c_i) \cap V(S_1)$ is not empty. We depict P in the above picture with a dotted line. In particular, P has length at least one, i.e., $c_j \notin N_{C,H}(c_i)$ by Lemma 8.2. Hence we obtain indeed a K_4-minor. Consequently, no vertex of $N_{C,H}(c_i)$

is contained in A, and thus $N_{C,H}(c_i) \subseteq B$.

It remains to show that H is a (1,2)-supergraph of $H[B]$: It is immediate that an edge e between two vertices in B is present in H if and only if it is present in $H[B]$. By the definition of B, H and $H[B]$ cannot have a different number of 2-paths between two different vertices b_1 and b_2 in B, unless $\{b_1, b_2\} = \{c_{i-1}, c_{i+1}\}$. However, regarding the latter case, all common neighbours of c_{i-1} and c_{i+1} are contained in $N_{C,H}(c_i) \subseteq B$ and thus the claim also holds for those two vertices. □

Corollary 8.4. *Let H be a biconnected K_4-minor-free graph containing an induced cycle $C = (c_0, \dots, c_{q-1}, c_0)$. If $q > 4$ then, for all $i \in \{0, \dots, q-1\}$, we have that at least one of $N_{C,H}(c_i)$ and $N_{C,H}(c_{i+1})$ has cardinality 1.*

Proof. Assume for contradiction that for some i, both, $N_{C,H}(c_i)$ and $N_{C,H}(c_{i+1})$, have cardinality greater than 1. We invoke Lemma 8.3 for C and i, which yields a separation (A, B) of H such that $C \setminus \{c_i\} \subseteq A$, $N_{C,H}(c_i) \subseteq B$, and $A \cap B = \{c_{i-1}, c_{i+1}\}$. However, by assumption, there exists $c' \in N_{C,H}(c_{i+1}) \setminus \{c_{i+1}\}$. Note further, that $c' \neq c_{i-1}$ as $q > 4$ and C is induced. Thus there is a path connecting $c_i \in B$ and $c_{i+2} \in A$ which does not pass through either one of c_{i-1} and c_{i+1} contradicting the assumption that (A, B) is a separation with $A \cap B = \{c_{i-1}, c_{i+1}\}$. □

Corollary 8.5. *Let H be a biconnected K_4-minor-free graph containing an induced cycle $C = (c_0, \dots, c_{q-1}, c_0)$. If $q \neq 4$ and H does not have a hardness gadget then, for all $i \in \{0, \dots, q-1\}$, we have that at least one of $N_{C,H}(c_i)$ and $N_{C,H}(c_{i+1})$ has cardinality 1.*

Proof. If $q > 4$ the statement follows from Corollary 8.4. If $q = 3$ then C is a triangle and the statement follows directly from Lemma 6.9. □

8.2. Pre-Hardness Gadgets and Obstructions.

Definition 8.6 (obstruction). Let B be a K_4-minor-free biconnected graph and let C be an induced cycle of B whose length is not 4. We say that B is an *obstruction with cycle C* if every even-cardinality walk-neighbour-set of C in B only contains vertices whose degree in B is 2. We say that B is an obstruction if, for some C, it is an obstruction with cycle C. We use $\mathrm{Cy}(B)$ to denote the set $\{C \mid B \text{ is an obstruction with cycle } C\}$.

Definition 8.7 (pre-hardness gadget). Let J be a connected graph. We say that J is a *pre-hardness gadget* if, for every (1,2)-supergraph H of J without K_4-minors, H has a hardness gadget.

Note that if J is a biconnected graph that is a pre-hardness gadget, then every K_4-minor-free graph H which contains J as a biconnected component has a hardness gadget.

It will be convenient to establish the following special case of an obstruction.

Lemma 8.8. *Let J be a K_4-minor-free biconnected graph such that the largest induced cycle of J is a square. If J contains a triangle then J is either a pre-hardness gadget or an obstruction.*

Proof. Let (a, b, c, a) be a triangle of J, let $a_1, \dots, a_k$ be the common neighbours of b and c with $a = a_1$, let $b_1, \dots, b_\ell$ be the common neighbours of a and c with $b = b_1$, and let $c_1, \dots, c_m$ be the common neighbours of a and b with $c = c_1$. If at least two of k, ℓ, and m are at least 2, then J is a strong hardness gadget by Lemma 6.9. In particular, every strong hardness gadget is also a pre-hardness gadget. If $k = \ell = m = 1$ then J is a pre-hardness gadget by Corollary 5.7, as follows. Let

H be a (1,2)-supergraph of J, let $q = 3$, and let $C = (a, b, c, a)$. Then $|N_{C,H}(a)| = |N_{C,H}(b)| = |N_{C,H}(c)| = 1$ since $k = \ell = m = 1$. Also, suppose for contradiction that there exists a walk $D = (d_a, d_b, d_c, d_a)$ with $d_a \in \Gamma_H(a) \setminus \{b, c\}$, $d_b \in \Gamma_H(b) \setminus \{a, c\}$ and $d_c \in \Gamma_H(c) \setminus \{a, b\}$. Consequently, as we do not allow self-loops in H, $d_a \neq a$, $d_b \neq b$ and $d_c \neq c$. Then the vertices d_a, a, b, c induce a K_4-minor (contract the edges $\{b, d_b\}$ and $\{c, d_c\}$ to obtain a K_4).

Hence assume w.l.o.g. that $k > 1$ and $\ell = m = 1$. If k is odd then J is a pre-hardness gadget by Lemma 6.10. If k is even and all a_j have degree 2 then J is an obstruction. Otherwise, for some $j \in \{1, \ldots, k\}$, let a_j have degree at least 3. As J is biconnected, there exists a shortest (induced) path P of length at least 2 from a_j to one of the vertices b, c or to some a_i with $i \in [k] \setminus \{j\}$. The internal vertices of P are disjoint from b, c and $\{a_i \mid i \in [k]\}$. If the endpoint of P is one of the other a_i, we obtain a K_4-minor, hence the endpoint must be b or c; suppose w.l.o.g. that it is c. As the largest induced cycle of J is a square, P has either length 2 or 3. In the former case, we obtain a strong (and thus also a pre-) hardness gadget by Lemma 6.9. In the latter case, J is a strong (and thus also a pre-) hardness gadget by Lemma 6.11. ∎

LEMMA 8.9. *Let H be a biconnected K_4-minor-free graph. If H contains an induced cycle of length at least 5 then H is either an obstruction or a pre-hardness gadget.*

Proof. We perform induction on $|V(H)|$: Let $C = (c_0, \ldots, c_{q-1}, c_0)$ be an induced cycle of length $q \geq 5$. If H is not an obstruction then, by Definition 8.6 there exists i such that $N_{C,H}(c_i)$ has even cardinality and contains a vertex of degree not equal to 2. Assume w.l.o.g. that $i = 1$. So we can assume that $N_{C,H}(c_1) = \{c_1^1, \ldots, c_1^k\}$ where $k > 0$ is even and $\deg_H(c_1^1) \neq 2$.

We invoke Lemma 8.3 and obtain a separation (A, B) of H such that $C \setminus \{c_1\} \subseteq A$, $N_{C,H}(c_1) \subseteq B$ and $A \cap B = \{c_0, c_2\}$. Furthermore, H is a (1,2)-supergraph of $H[B]$. Now consider the neighbours of c_1^1: We have that $c_0 \in \Gamma_H(c_1^1)$ and $c_2 \in \Gamma_H(c_1^1)$ by the definition of $N_{C,H}(c_1)$. As $\deg_H(c_1^1) \neq 2$, there exists another neighbour $w \in \Gamma_H(c_1^1)$. By the properties of the separation (A, B), for any $w \in \Gamma_H(c_1^1) \setminus \{c_0, c_2\}$, $w \in B$.

Claim A: *There is a vertex w in $\Gamma_H(c_1^1) \setminus \{c_0, c_2\}$ and an induced path P in $H[B]$ from w to either c_0 or c_2 such that all internal vertices of P are contained in $B \setminus (N_{C,H}(c_1) \cup \{c_0, c_2\})$. Furthermore, no internal vertex of P is a neighbour of c_1^1.*
Proof: Let $w' \in \Gamma_H(c_1^1) \setminus \{c_0, c_2\}$. As H is biconnected, the vertex c_1^1 is not an articulation point. Consequently, there exists a path P' from w' to c_0 not containing c_1^1 as internal vertex. We can assume P' to be induced by taking possible "shortcuts". W.l.o.g. we have that P' does not visit c_2 as internal vertex as, otherwise, we can just continue with c_2 instead of c_0.

Assume first that P' contains a vertex in $A \setminus B$. As (A, B) is a separation and $w' \in B$, we have that P' is of the form

$$w' \xrightarrow{P_1} x \xrightarrow{P_2} c_0\,,$$

such that P_1 is contained in $H[B]$ and $x \in A \cap B = \{c_0, c_2\}$. However, as P' is a path that does not contain c_2 as internal vertex, we obtain that $P_2 = \emptyset$ and $x = c_0$, contradicting the assumption.

Next assume that P' contains an internal vertex z in $N_{C,H}(c_1) \setminus \{c_1^1\}$; we obtain the contradiction by identifying a K_4-minor in H as depicted in Figure 8. We have now shown that there is an induced path P' in $H[B]$ from w' to c_0 or c_2 such that all internal vertices of P' are contained in $B \setminus (N_{C,H}(c_1) \cup \{c_0, c_2\})$. Now choose w

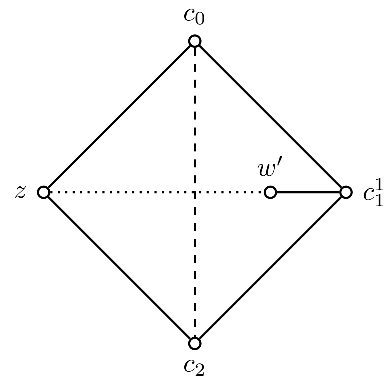

Fig. 8: The K_4-minor used in the proof of Claim A in Lemma 8.9. The dashed line depicts the remainder of the cycle, and the dotted line depicts P'.

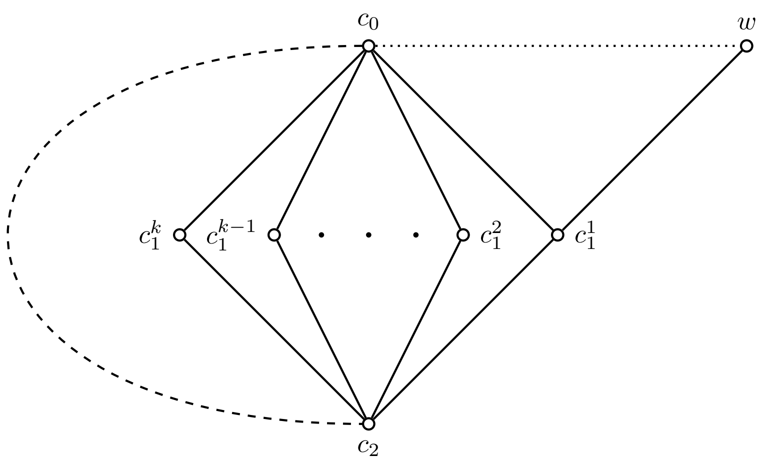

Fig. 9: Illustration of cycle C consisting of the dashed line and one of the vertices c_1^i, and path P (dotted) in the proof of Lemma 8.9.

to be the first neighbour of c_1^1 along P' from c_0 or c_2, respectively, and let P be the sub-path of P' going from c_0 or c_2, respectively, to w.

■

We assume in the remainder of the proof that the Claim A holds for c_0; the case of c_2 is completely symmetric (by substituting every subsequent appearance of c_0 by c_2 and vice versa). For convenience, we also provide an illustration of our current situation in Figure 9. For the remainder of the proof, we need the following observation:

Claim B: *$H[B]$ is biconnected.*
Proof: By Menger's Theorem, we have to show that there are two internally vertex-disjoint paths (in $H[B]$) between every pair of different vertices x and y in B. As H is

biconnected, there are two such paths Q_1 and Q_2 connecting x and y in H. If $\{x, y\} = \{c_0, c_2\}$, then the claim follows immediately as $N_{C,H}(c_1) \subseteq B$ and $|N_{C,H}(c_i)| \geq 2$.

Hence we can assume that $\{x, y\} \neq \{c_0, c_2\}$. The next step is to show that at least one of Q_1 and Q_2 is fully contained in $H[B]$. If $\{x, y\}$ intersects $\{c_0, c_2\}$ (for example, if $y = c_2$) then this is clear because c_0 can only be on one of Q_1, Q_2. Otherwise, suppose that one of the paths, say Q_2, from x to y, leaves $H[B]$. It leaves by one of the vertices in the separator $\{c_0, c_2\}$ and returns by the other. So Q_1 stays within $H[B]$. If Q_2 is also fully contained in $H[B]$ we are done.

Otherwise we have that w.l.o.g. (otherwise switch c_0 and c_2 and proceed symmetrically):

$$Q_2 = x \overset{Q_2^1}{\to} c_0 \overset{Q_2^2}{\to} c_2 \overset{Q_2^3}{\to} y\,,$$

where Q_2^1 and Q_2^3 are in $H[B]$ and Q_2^2 is non-empty in $H[A \setminus B]$. Next we claim that Q_1 contains at most one vertex in $N_{C,H}(c_1)$ as internal vertex. Assuming otherwise, we have

$$Q_1 = x \overset{Q_1^1}{\to} c_1^1 \overset{Q_1^2}{\to} c_1^2 \overset{Q_1^3}{\to} y\,,$$

where $c_1^1 \neq c_1^2 \in N_{C,H}(c_1)$. As Q_1 is fully contained in $H[B]$, we obtain a K_4-minor, unless Q_1^2 contains c_0 or c_2 as internal vertices: The K_4-minor is induced by c_0, c_1^1, c_1^2, c_2 — note that c_0 is connected to c_2 by $C \setminus \{c_1\}$, and c_1^1 is connected to c_1^2 by Q_1^2.

Thus we can assume that Q_1^2 contains c_0 or c_2 as internal vertices.

- If c_2 is an internal vertex of Q_1^2, then $y \neq c_2$. In this case, however, Q_1 and Q_2 share c_2 as internal vertex, which leads to a contradiction.
- If c_0 is an internal vertex of Q_1^2, then $x \neq c_0$. In this case, however, Q_1 and Q_2 share c_0 as internal vertex, which leads to a contradiction.

Consequently, Q_1 contains at most one vertex in $N_{C,H}(c_1)$. As $N_{C,H}(c_1)$ is of even positive cardinality, there exists hence a vertex $z \in N_{C,H}(c_1)$ which is not part of Q_1. Finally, this enables us to modify Q_2 be substituting Q_2^2 by the path c_0, z, c_2. The resulting path is fully contained in $H[B]$ and, by the previous analysis, internally vertex-disjoint from Q_1. This concludes the proof of Claim B. ■

We proceed with the following claim.

Claim C: *If $H[B]$ is an obstruction, then so is H.*

Proof: If $H[B]$ is an obstruction then it contains an induced cycle D satisfying the requirements of Definition 8.6. For the sake of readability, we state those requirements explicitly: The graph $H[B]$ contains an induced cycle $D = (d_0, \ldots, d_{r-1}, d_0)$ for some $r \neq 4$. Furthermore, we have that for all i, every vertex in $N_{D,H[B]}(d_i)$ has degree 2 in $H[B]$, unless $|N_{D,H[B]}(d_i)|$ is odd.

We claim that H is an obstruction with cycle D: Observe that

$$N_{D,H[B]}(d_i) = \Gamma_{H[B]}(d_{i-1}) \cap \Gamma_{H[B]}(d_{i+1}) = \Gamma_H(d_{i-1}) \cap \Gamma_H(d_{i+1}) = N_{D,H}(d_i)\,,$$

where the second equality is true as H is a (1,2)-supergraph of $H[B]$. Consequently, it remains to show that for all i with $|N_{D,H}(d_i)|$ even, every vertex in $N_{D,H}(d_i)$ has degree 2 in H. For the sake of contradiction, we assume w.l.o.g. that $N_{D,H}(d_1)$ is of even cardinality and contains a vertex $\hat{d}_1$ such that $\hat{d}_1$ has degree 2 in $H[B]$, but degree at least 3 in H. As the separator of (A, B) is $\{c_0, c_2\}$, the only possibility for this to happen is $\hat{d}_1 = c_0$ or $\hat{d}_1 = c_2$. However, $\hat{d}_1 = c_0$ is impossible, as c_0 has at least three neighbours already in $H[B]$: c_0 is adjacent to every vertex in $N_{C,H}(c_1)$, which is of positive even cardinality (i.e., of size at least 2), and c_0 is adjacent to the first vertex in the path P from c_0 to w (see Figure 9).

Hence the remaining possibility is $\hat{d}_1 = c_2$. Recall that $\hat{d}_1 = c_2$ has neighbours $c_1^1, \ldots, c_1^k$ in B. So if there are only two of them, then $k = 2$ and $|N_{C,H}(c_1)| = 2$. However, as $\hat{d}_1 \in N_{D,H}(d_1)$ has degree 2 in $H[B]$, and c_1^1 and c_1^2 are adjacent to $\hat{d}_1 = c_2$ in $H[B]$, we obtain that $\{c_1^1, c_1^2\} = \{d_0, d_2\}$ — recall that $\hat{d}_1 \in \Gamma_{H[B]}(d_0) \cap \Gamma_{H[B]}(d_2)$. Finally, $N_{D,H}(d_1)$ has positive, even cardinality. Thus there exists a vertex $d_1' \neq \hat{d}_1$ in $N_{D,H}(d_1)$ which is also adjacent to d_0 and d_2. This yields the following K_4-minor of H; note that $N_{D,H}(d_1) \subseteq B$ and the dashed line is $C \setminus c_1$, which is in A.

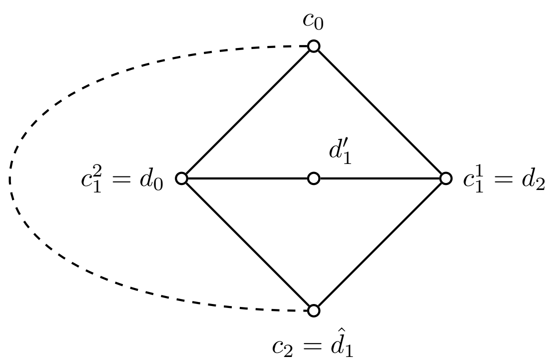

This concludes the proof of Claim C. ■

In what follows, we perform a case distinction along the length L of the largest induced cycle in $H[B]$.

(I) $L \geq 5$. This allows us to invoke the induction hypothesis to the graph $H[B]$; note that $|V(H[B])| = |B|$ is indeed strictly smaller than $|V(H)|$ as the cycle C has length at least 5, and thus A is not empty. Furthermore, $H[B]$ is biconnected by Claim B. If $H[B]$ is a pre-hardness gadget, then so is H, as H is a (1,2)-supergraph of $H[B]$. If $H[B]$ is an obstruction, then so is H by Claim C.

(II) $L \leq 4$. Consider again the path P in Figure 9. By the assumption of this case, P is either an edge or a 2-path. If P is an edge, then $H[B]$ satisfies all conditions of Lemma 8.8. Consequently, $H[B]$ is a pre-hardness gadget or an obstruction. In the former case, we are done as H is a (1,2)-supergraph of $H[B]$ and thus also a pre-hardness gadget. In the latter case, we obtain that H is an obstruction as well by invoking Claim C.
Finally, assume that P is a 2-path. We claim that H is a pre-hardness gadget. To this end, let H' be a K_4-minor-free (1,2)-supergraph of H. Then H' contains the following subgraph:

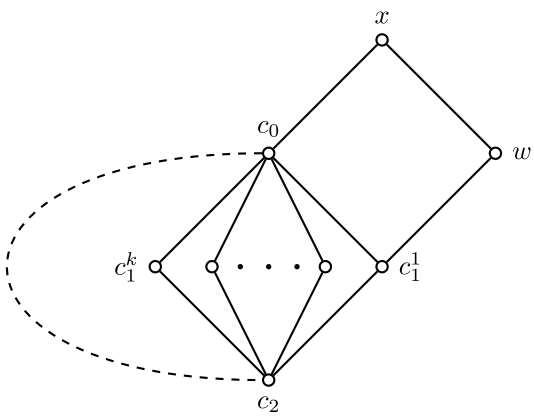

In particular, we have that c_1^1 and c_1^k have no common neighbours in H' apart from c_0 and c_2: This is due to the fact that they have no further common neighbours in H, as otherwise H has a K_4-minor similarly as in the proof of Claim A, and as H' is a (1,2)-supergraph, it cannot add common neighbours to vertices. Furthermore, we have that $k > 0$ is even and that $c_1^1, \ldots, c_1^k$ are all common neighbours of c_0 and c_2 in H and thus in H'. We apply Lemma 6.5 to the subgraph of H' induced by the vertices $c_1^k, c_0, x, w, c_1^1, c_2$ and obtain a hardness gadget in H', unless this subgraph, call it F, is of type V. By the previous argument, the only possibility for F being of type V is k being strictly greater than 2. However, as k is even, we obtain a hardness gadget in H' in this case as well: We found an instance of Lemma 6.7. ∎

8.3. K_4-minor-free Component Lemma.

LEMMA 8.10 (K_4-minor-free Component Lemma). *Let B be a biconnected K_4-minor-free graph. If B is not an edge then at least one of the following is true:*
(a) B is a diamond.
(b) B is an obstruction.
(c) B is an impasse.
(d) For every K_4-minor-free graph H containing B as a biconnected component, H has a hardness gadget.

Proof. Let L be the size of the largest induced cycle of B. Note that $L \geq 3$ is well-defined as B is biconnected, but not an edge. If $L \geq 5$ we obtain by Lemma 8.9 that B is either a pre-hardness gadget or an obstruction. In the latter case, (b) holds. In the former case, (d) holds, as every K_4-minor-free graph H containing B as a biconnected component is a (1,2)-supergraph of B.

If $L \leq 4$ and B contains a triangle, then B is either a pre-hardness gadget or an obstruction by Lemma 8.8. Similarly as before, (b) or (d) hold.

In the remaining case, B is chordal bipartite and we can invoke Lemma 6.17, yielding that either (a), (c) or (d) hold. ∎

9. K_4-minor-free Graphs.

9.1. Suitable Connectors.

Definition 9.1 (suitable connector). Let H be a graph, let B be a biconnected component of H, and let $A \subseteq V(B)$ be a set of articulation points of H. We say that (B, A) is a *suitable connector* in H if one of the following cases holds:

- B is an edge $\{a,b\}$ and $A=\{a,b\}$, or
- B is a diamond (Definition 6.16) that contains an edge $\{a,b\}$ such that $A=\{a,b\}$, or
- B is an impasse (Definition 6.15) that has a pair of connectors (a,b) such that $A=\{a,b\}$, or
- B is an obstruction (Definition 8.6). In this case (B,A) is a suitable connector in H if there is a cycle $C\in\mathrm{Cy}(B)$ such that $A=\{c\in C\mid |N_{C,H}(c)| \text{ is even}\}$. Note that A could be the empty set. If (B,A) is a suitable connector in H then we fix a particular cycle $C(B,A)\in\mathrm{Cy}(B)$ such that

$$A=\{c\in C(B,A)\mid \text{the cardinality of } N_{C(B,A),H}(c) \text{ is even}\}.$$

(It does not matter if there are multiple possibilities for $C(B,A)$ in $\mathrm{Cy}(B)$ — we just fix one, for example, the lexicographically least one.)

LEMMA 9.2. *Let B be a biconnected component in an involution-free graph H. If B is a diamond then there exists a set $A\subseteq V(B)$ of articulation points of H such that (B,A) is a suitable connector in H.*

Proof. If B is a diamond with vertices as given in Definition 6.16, then as H is involution-free there exist articulation points $a\in\{s,t\}$ and $b\in\{x_1,\ldots,x_k\}$. Hence, for $A=\{a,b\}$, (B,A) is a suitable connector in H. ◻

LEMMA 9.3. *Let B be a biconnected component in an involution-free graph H. If B is an impasse then there exists a set $A\subseteq V(B)$ of articulation points of H such that (B,A) is a suitable connector in H.*

Proof. If B is an impasse with vertices as given in Definition 6.15, then as H is involution-free there exist articulation points $a\in\{v_1,y_1,\ldots,y_k\}$ and $b\in\{v_3,z_1,\ldots,z_\ell\}$. Note that (a,b) is a pair of connectors (cf. Definition 6.15) and hence, for $A=\{a,b\}$, (B,A) is a suitable connector in H. ◻

LEMMA 9.4. *Let B be a biconnected component in an involution-free graph H. If B is an obstruction then there exists a set $A\subseteq V(B)$ of articulation points of H such that (B,A) is a suitable connector in H.*

Proof. If B is an obstruction then there exists a cycle C with $C\in\mathrm{Cy}(B)$. Let $c\in C$ such that $|N_{C,H}(c)|$ is even. By definition of an obstruction, every vertex in $|N_{C,H}(c)|$ has degree 2 in B. Since $c\in N_{C,H}(c)$, $|N_{C,H}(c)|\geq 2$. Therefore, as H is involution-free, at least one vertex in $N_{C,H}(c)$ is an articulation point of H. By renaming vertices, we can assume without loss of generality that c is an articulation point. Hence, for $A=\{c\in C\mid \text{the cardinality of } N_{C,H}(c) \text{ is even}\}$, (B,A) is a suitable connector in H, where $C(B,A)=C$. ◻

9.2. Finding a Suitable Subtree. In this section we will use the notion of rooted trees. Given a tree T and a vertex r in T, (T,r) is a *rooted tree* and the *tree-order* $<_r$ induced by r (on T) is the partial order of the vertices of T, where for vertices u and v of T we have $u<_r v$ if and only if the unique path from r (the root) to v passes through u. Such a partial order gives rise to the standard notion of child, parent, ancestor and descendant. In order to clarify which tree-order we are referring to we speak of an r-child, r-parent, r-ancestor and r-descendant when we mean child, parent, ancestor and descendant with respect to $<_r$.

For a connected graph H, recall the definition of the block-cut tree $\mathsf{BC}(H)$ from Definition 4.3.

Definition 9.5 (R-open, R-closed). Let H be a connected graph, let a be a cut vertex in $\mathsf{BC}(H)$, and let R be a block in $\mathsf{BC}(H)$. If a has exactly one descendant with respect to $<_R$ in $\mathsf{BC}(H)$ and this descendant is a block in $\mathsf{BC}(H)$ that is an edge, then a is *R-closed* (in $\mathsf{BC}(H)$). Otherwise, a is *R-open* (in $\mathsf{BC}(H)$).

Definition 9.6 (suitable subtree, closed). Let H be a connected graph. Let T be a subtree of $\mathsf{BC}(H)$. We say that T is *suitable* if it has the following properties:

1. For every block B in T, $(B, \Gamma_T(B))$ is a suitable connector in H (Definition 9.1).
2. Every cut vertex of T has degree at most 2 in T.

A suitable subtree T is *closed* if there exists a block R in T such that every cut vertex that is a leaf in T is R-closed in $\mathsf{BC}(H)$.

Lemma 9.7. *Let H be a connected graph and let T be a suitable subtree of $\mathsf{BC}(H)$. Let R and R' be distinct blocks in T and let a be a cut vertex that is a leaf in T. If a is R-closed in $\mathsf{BC}(H)$ then it is R'-closed in $\mathsf{BC}(H)$.*

Proof. Let B be a block of $\mathsf{BC}(H)$. We show that B is an R-descendant of a in $\mathsf{BC}(H)$ if and only if it is an R'-descendant. From this it follows immediately that if a is R-closed in $\mathsf{BC}(H)$ then it is R'-closed in $\mathsf{BC}(H)$. Let B be an R-descendant of a. Since a is a leaf of T and R is in T it follows that B is not in T. Since R, R' and a are all in T, there is a path in T from R' to a and consequently this path does not contain B. Hence the unique path from R' to B goes through a, which means that B is an R'-descendant of a in $\mathsf{BC}(H)$. It is analogous to show that if B is an R'-descendant of a it is also an R-descendant. ∎

The following lemma gives the initialisation for finding a closed suitable subtree (which is then done in Lemma 9.9).

Lemma 9.8. *Let H be an involution-free, connected graph such that every biconnected component of H is an edge, a diamond, an impasse or an obstruction. Then there exists a biconnected component B_0 and a set of articulation points $A_0 \subseteq V(B_0)$ such that (B_0, A_0) is a suitable connector in H and hence $T(B_0) = (\{B_0\} \cup A_0, \{\{B_0, a\} \mid a \in A_0\})$ is a suitable subtree of $\mathsf{BC}(H)$.*

Proof. First note that if all biconnected components of H are edges, then there is at least one edge between articulation points as H is involution-free and therefore H is not a star. Therefore, H contains a biconnected component R that is one of the following: a diamond, an impasse, an obstruction, or an edge for which both endpoints are articulation points of H. In the first three cases we can use Lemmas 9.2, 9.3 or 9.4, respectively, to obtain a suitable connector. If B_0 is an edge $\{a, b\}$ where both end points are articulation points, then $(B_0, \{a, b\})$ is a suitable connector. Then it is immediate that $T(B_0)$ is a suitable subtree of $\mathsf{BC}(H)$. ∎

Lemma 9.9. *Let H be an involution-free, connected graph such that every biconnected component of H is an edge, a diamond, an impasse or an obstruction. Then there exists a closed suitable subtree of $\mathsf{BC}(H)$.*

Proof. Let $B_0, A_0, T(B_0)$ be as given by Lemma 9.8. Algorithm 9.1 keeps track of a suitable subtree T of $\mathsf{BC}(H)$, a block R of T, and the set $A(T)$ of leaves of T that are cut vertices (i.e., that are articulation points of H).

Algorithm 9.1

$T \leftarrow T(B_0)$
$R \leftarrow B_0$
$A(T) \leftarrow A_0$
while $A(T)$ contains an R-open cut vertex a^*
    // Invariant: All elements of $A(T)$ are B_0-descendants of R.
    if there is a suitable connector (B, A) in H such that B is an R-child of a^* and $a^* \in A$
        // By the invariant, every element of $A \setminus \{a^*\}$ is a B_0-descendant of a^*.
        $V \leftarrow V(T) \cup \{B\} \cup A$
        $E \leftarrow E(T) \cup \{\{B, a\} \mid a \in A\}$
        $T \leftarrow (V, E)$
        $A(T) \leftarrow (A(T) \cup A) \setminus \{a^*\}$
    else
        Choose a suitable connector (B, A) in H such that B is an R-child of a^* in $\mathsf{BC}(H)$.
        // By the invariant, every element of A is a B_0-descendant of a^*.
        $V \leftarrow \{B\} \cup A$
        $E \leftarrow \{\{B, a\} \mid a \in A\}$
        $T \leftarrow (V, E)$
        $R \leftarrow B$
        $A(T) \leftarrow A$

We now show that Algorithm 9.1 is well-defined and finds a closed suitable subtree.[6] In order to show that the algorithm is well-defined note that any R-open cut vertex a^* is an articulation point of H and therefore is adjacent to at least two blocks of $\mathsf{BC}(H)$. At most one of these blocks can be an R-parent. Therefore a^* has an R-child in $\mathsf{BC}(H)$. If there is such an R-child B that is a diamond, an impasse, or an obstruction, then by Lemmas 9.2, 9.3 or 9.4, respectively, there exists a suitable connector of the form (B, A). If otherwise all R-children of a^* are edges then a^* has at least one such R-child $B = \{a^*, b\}$ for which b is an articulation point (as a^* is R-open and H is involution-free). Therefore $(B, \{a^*, b\})$ is a suitable connector. Thus, the algorithm is well-defined as we can always choose a suitable connector (B, A) where B is an R-child of a^*.

We next show that at any point during the algorithm, T is a suitable subtree of $\mathsf{BC}(H)$, R is a block in T, and $A(T)$ is the set of leaves of T that are cut vertices of $\mathsf{BC}(H)$. First note that in the initialisation this clearly holds by Lemma 9.8. We show that after each update these properties still hold. Note that if we update T, R, and $A(T)$ as part of the else-block then $R = B$ is the only block in T, $\Gamma_T(B) = A$, and (B, A) is a suitable connector. Thus, T is a suitable subtree. Furthermore, the cut vertex leaves of T are precisely the elements of A and we have $A(T) = A$, as required.

If otherwise we update T and $A(T)$ as part of the if-block then

1. The block R continues to be a vertex of T.
2. We add precisely one block B together with the articulation points A and the edges $\{\{B, a\} \mid a \in A\}$, which ensures that $\Gamma_T(B) = A$ and hence $(B, \Gamma_T(B))$

[6] Since the graph H is fixed, the running time of Algorithm 9.1 is not important for us. What is important is that the algorithm gives us a (constructive) proof that such a closed suitable subtree exists.

is a suitable connector.

3. All cut vertices in $A \setminus \{a^*\}$ are leaves in T and since a^* was a leaf before the update, it now has degree 2 in T.

Consequently, T is a suitable subtree after the update. Furthermore, we remove a^* from $A(T)$ as it now has degree 2 in T, and we add the cut vertices $A \setminus \{a^*\}$ to $A(T)$ since they are leaves in T.

We have established that at any point during the algorithm, T is a suitable subtree of $\mathsf{BC}(H)$, R is a block in T, and $A(T)$ is the set of leaves of T that are cut vertices. It remains to show that Algorithm 9.1 terminates (in which case it is immediate that T is a closed suitable subtree). Note that with each iteration we remove a vertex a^* from $A(T)$. With each iteration we may also add some vertices to $A(T)$. As noted in Algorithm 9.1, the vertices that are added in each iteration are always B_0-descendants in $\mathsf{BC}(H)$ of the vertex a^* that is deleted. It follows immediately that Algorithm 9.1 terminates as we only consider finite graphs. ∎

9.3. Suitable Subtrees without Obstructions.

LEMMA 9.10. *Let H be a connected graph and let T be a closed suitable subtree of $\mathsf{BC}(H)$. If no block of T is an obstruction then H has a hardness gadget.*

Proof. As T does not contain an obstruction, the degree of every block in T is 2. Together with the fact that every cut vertex has degree at most 2, this implies that, for a non-negative integer q, T is a path of the form $(b_0, B_1, b_1, B_2, \dots, B_q, b_q)$, where $B_1, \dots, B_q$ are blocks, i.e. biconnected components of H, and $b_0, \dots, b_q$ are cut vertices, i.e. articulation points of H. Since T is closed it contains at least one block R and therefore $q \geq 1$. Furthermore, for each $i \in [q]$, $(B_i, \{b_{i-1}, b_i\})$ is a suitable connector. And since B_i is no obstruction, one of the following holds:

- B_i is an edge $\{b_{i-1}, b_i\}$, or
- B_i is a diamond that contains the edge $\{b_{i-1}, b_i\}$, or
- B_i is an impasse such that (b_{i-1}, b_i) is a pair of connectors.

Since T is closed, there is a block R among $B_1, \dots, B_q$ such that both b_0 and b_q are R-closed. By Lemma 9.7, b_0 is B_1-closed and b_q is B_q-closed. It follows that $|\Gamma_H(b_0) \setminus V(B_1)| = 1$ and $|\Gamma_H(b_q) \setminus V(B_q)| = 1$.

Thus, we can apply Lemma 7.14 to obtain that H has a hardness gadget or otherwise there exists $L_q \subseteq \Gamma_{B_q}(b_q)$ such that (L_q, b_q) is a good start in B_q. Since $\Gamma_H(b_q) \setminus V(B_q)$ has odd cardinality, this means that (L_q, b_q) is a good stop in B_q. Then Lemma 7.14 ensures that H has a hardness gadget in this case as well. ∎

9.4. Suitable Subtrees with Obstructions. The goal of this section is to prove Lemma 9.24, which gives a hardness gadget in a connected K_4-minor-free graph using a closed suitable subtree that contains an obstruction. In order to find this hardness gadget we use Lemma 5.6, which derives a hardness gadget based on the generalised cycle gadget from Definition 5.4. The sets of vertices $\mathcal{C}_0, \dots, \mathcal{C}_{q-1}$ from Lemma 5.6 will correspond to the walk-neighbour-sets of a specific closed walk W. With Algorithms 9.2 and 9.3 we define this walk W — it is the output of Algorithm 9.3. In Lemmas 9.17, 9.18 and 9.19 we establish that the algorithms are well-defined and give as output a closed walk in H whose length is at least 3, and not equal to 4. In Figure 10 we give an example that illustrates how W is derived. In Lemmas 9.22 and 9.23 we then show that the walk-neighbour-sets of W satisfy certain properties required to apply Lemma 5.6. In the proof of Lemma 9.24 we put all the pieces together and establish the remaining necessary properties of W.

Definition 9.11 (obstruction-free path, proper). Let H be a connected graph and let T be a closed suitable subtree of $\mathsf{BC}(H)$. A path in T is *obstruction-free* if it does not contain a block that is an obstruction. An obstruction-free path is *proper* if its endpoints are cut vertices of $\mathsf{BC}(H)$. Note that it is possible that a proper obstruction-free path has length 0. Then it is of the form (v) where v is a cut vertex of $\mathsf{BC}(H)$.

Definition 9.12 ($\mathrm{P}_H(a,b)$). Let H be a graph and let a and b be vertices of H. If $a = b$ then $\mathrm{P}_H(a,b) = (a)$. If $a \neq b$ then $\mathrm{P}_H(a,b)$ is a shortest path from a to b in H.

In Definition 9.12, it is of course possible that H might have multiple shortest paths from a to b. In this case, it doesn't matter which of these is chosen to be $\mathrm{P}_H(a,b)$ — for concreteness, the reader may assume that $\mathrm{P}_H(a,b)$ is the lexicographically least of these. (In fact, when we use the definition, this shortest path will turn out to be unique.)

Lemma 9.13. *Let H be a connected graph and let T be a closed suitable subtree of $\mathsf{BC}(H)$. For a non-negative integer q, let $P = (b_0, B_1, b_1, B_2, \dots, B_q, b_q)$ be a proper obstruction-free path in T. Then $\mathrm{P}_H(b_0, b_q)$ is the unique shortest path from a to b in H. It passes through $b_0, b_1, \dots, b_q$ in order. For $i \in [q]$, the subpath of $\mathrm{P}_H(b_0, b_q)$ that connects b_{i-1} and b_i is either an edge or it is of the form (b_{i-1}, v, b_i), where v is the unique common neighbour of b_{i-1} and b_i in H.*

Proof. Since P is a proper obstruction-free path, $b_0, b_1 \dots, b_q$ are cut vertices and $B_1, \dots, B_q$ are blocks. Since T is a suitable subtree and P is obstruction free, for each $i \in [q]$, $(B_i, \{b_{i-1}, b_i\})$ is a suitable connector, where B_i is an edge, diamond or impasse. Since $B_1, \dots, B_q$ are biconnected components, every path from b_0 to b_q traverses $b_0, b_1 \dots, b_q$ in order. The shortest path from b_0 to b_q is unique, if for each $i \in [k]$, the shortest path from b_{i-1} to b_i is unique. If B_i is an edge or diamond, this is clearly the case since then the shortest path from b_{i-1} to b_i is an edge. If B_i is an impasse then (b_{i-1}, b_i) is a pair of connectors of B_i and by Definition 6.15 there is no edge between b_{i-1} and b_i, but there is a unique common neighbour v of b_{i-1} and b_i in H and consequently the unique shortest path from b_{i-1} to b_i is of the form (b_{i-1}, v, b_i), as required. ∎

Definition 9.14 (attachment point, exit, destination). Let H be a connected graph and let T be a closed suitable subtree of $\mathsf{BC}(H)$. Let a be a cut vertex that has an obstruction B as a neighbour in T. Then, since every cut vertex of T has degree at most 2, there is a unique maximal-length proper obstruction-free path P^* in T starting at a. Let b be the other endpoint of P^* (possibly $P^* = (a)$ in which case $b = a$). The vertex a is an *attachment point of* (T, B) if b is a leaf in T. Otherwise, a is an *exit of* (T, B). In this case, b is adjacent to a block $B' \neq B$ which is an obstruction. We say that (b, B') is the *destination* of a in T.

At the beginning of this section we outlined our plan to define a particular closed walk W. We chose the names in Definition 9.14 since W will *exit* an obstruction when it encounters an exit, and it will then proceed towards the destination of that exit. The walk W will not exit an obstruction when it encounters an attachment point. However, W will be designed so that every even-cardinality walk-neighbour-set of W contains an attachment point, and the structure that is *attached* to such a point will allow us to construct a hardness gadget.

Definition 9.15 (concatenation "+"). Let $W = (w_0, \dots, w_k)$ and $W' = (w'_0, \dots, w'_\ell)$ be two walks such that $w_k = w'_0$. If $k = 0$ then the *concatenation* $W + W'$ of

W with W' is equal to W'. Similarly, if $\ell = 0$, it is equal to W. If both k and ℓ are positive then $W + W' = (w_0, \ldots, w_{k-1}, w_k, w'_1 \ldots, w'_\ell)$.

Definition 9.16 ($D(C)$, $\mathrm{W}_C(a)$, $\mathrm{W}_C(a,b)$). For an integer $q \geq 3$, let $C = (c_0, \ldots, c_{q-1}, c_0)$ be a cycle in a graph H. Then $D(C)$ is the cyclic order induced by the order in which the walk C traverses the vertices $\{c_0, \ldots, c_{q-1}\}$. For $a \in C$, $\mathrm{W}_C(a)$ is the walk from a to itself following all of the vertices of C in the order given by $D(C)$. For $a, b \in C$, $\mathrm{W}_C(a,b)$ is the walk from a to b along C in the order given by $D(C)$.

Algorithm 9.2 $\textsc{ExitWalk}(T, a^*, B, \ell, a_0)$

Input: A closed suitable subtree T of $\mathsf{BC}(H)$ of a connected graph H, a cut vertex a^* in T, an obstruction B that is a block in T such that $\mathsf{dist}_T(a^*, B) = \ell$, and an exit a_0 of (T, B)

$C \leftarrow C(B, \Gamma_T(B))$
$\{a_0, \ldots, a_k\} \leftarrow$ The exits of (T, B) in the order of $D(C)$, starting from a_0
if $k = 0$
  $W \leftarrow \mathrm{W}_C(a_0)$.
else
  $\{(b_1, B_1), \ldots, (b_k, B_k)\} \leftarrow$ The destinations of $a_1, \ldots, a_k$, respectively
  $W \leftarrow \mathrm{W}_C(a_0, a_1)$
  for $i = 1, \ldots, k$
    $r_i \leftarrow \mathsf{dist}_T(B, b_i)$
    $W \leftarrow W + \mathrm{P}_H(a_i, b_i) + \textsc{ExitWalk}(T, a^*, B_i, \ell + r_i + 1, b_i) + \mathrm{P}_H(b_i, a_i) +$
    $\mathrm{W}_C(a_i, a_{i+1 \mod k+1})$

Output: W

Algorithm 9.3 $\textsc{Walk}(T, B')$

Input: A closed suitable subtree T of $\mathsf{BC}(H)$ of a connected graph H, an obstruction B' that is a block in T

if there is an exit a^* of (T, B')
  $(b^*, B^*) \leftarrow$ The destination of a^*
  $r^* \leftarrow \mathsf{dist}_T(a^*, b^*)$
  $W \leftarrow \textsc{ExitWalk}(T, a^*, B', 1, a^*) + \mathrm{P}_H(a^*, b^*) +$
    $\textsc{ExitWalk}(T, a^*, B^*, r^* + 1, b^*) + \mathrm{P}_H(b^*, a^*)$
else
  $W \leftarrow C(B', \Gamma_T(B'))$

Output: W

In Figure 10 we provide some illustrations of a graph H, a closed suitable subtree $T \in \mathsf{BC}(H)$, and the walk W returned by Algorithm 9.3. In order to gain intuition, it is probably useful to simulate $\textsc{Walk}(T, O_1)$. The exit a^* can be chosen to be a_2 with destination $(b^*, B^*) = (b_2, O_2)$. The variable r^* is set to 0. So the first part of W is the output of the call $\textsc{ExitWalk}(T, a_2, O_1, 1, a_2)$.

Let's start by considering that call. $\Gamma_T(O_1) = \{t_1, a_1, a_2\}$ and C is the cycle around O_1 shown in red. The exits are $\{a_2, a_1\}$ so the output W of this call starts by following the red cycle clockwise from a_2 to a_1. In the else-clause we have $k = 1$

and the destination of a_1 is (b_1, O_3). The walk next takes the unique shortest path from a_1 to b_1. Then there is a call to $\textsc{ExitWalk}(T, a_2, O_3, \ell, b_1)$, for some value of ℓ (the value of ℓ doesn't matter — it is just for accounting). The only exit of (T, O_3) is b_1, so this call returns a walk around the red cycle in O_3 from b_1 to itself. Finally, the call to $\textsc{ExitWalk}(T, a_2, O_1, 1, a_2)$ takes the unique shortest path back from b_1 to a_1 and finishes the red cycle in O_1 clock-wise, back to a_2. Thus, the output of $\textsc{ExitWalk}(T, a_2, O_1, 1, a_2)$ is a closed walk from a_2 to itself that covers all of the red edges in the picture, apart from the triangle in O_2.

This is concatenated with $P_H(a_2, b_2) = (a_2)$ (which does nothing). Then it is concatenated with the output of a call to $\textsc{ExitWalk}(T, a_2, O_2, 1, a_2)$. Now $(O_2, \{a_2\})$ is a suitable connector in H with $C(O_2, \{a_2\})$ equal to the red triangle in O_2, so C is assigned to be this triangle. The cut-vertex a_2 is the only exit of C, so this call returns the walk from a_2 to itself around C. Concatenating $P_H(a_2, b_2)$ with this does not change the output. The entire walk is coloured in red.

We now proceed to establish the correctness of Algorithms 9.2 and 9.3, and to prove some properties of the walks that they output.

Lemma 9.17. *All calls to* $\textsc{ExitWalk}(\cdot)$ *in Algorithms 9.2 and 9.3 have arguments that are feasible inputs to Algorithm 9.2.*

Proof. First, consider Algorithm 9.2, where for $i = 1, \dots, k$ we make a call $\textsc{ExitWalk}(T, a^*, B_i, \ell + r_i + 1, b_i)$. Observe that b_i is an exit of (T, B_i) by the definition of a destination (Definition 9.14). It remains to check that $\mathsf{dist}_T(a^*, B_i) = \ell + r_i + 1$. This is true since $\ell = \mathsf{dist}_T(a^*, B)$ and $r_i = \mathsf{dist}_T(B, b_i)$ using the fact that the (unique) path from a^* to B_i in the tree T goes from a^* to B then from B to b_i and then from b_i to B_i, where B_i is adjacent to b_i.

Second, consider Algorithm 9.3. The if-block makes two calls to $\textsc{ExitWalk}(\cdot)$ — one is $\textsc{ExitWalk}(T, a^*, B', 1, a^*)$ and the other is $\textsc{ExitWalk}(T, a^*, B^*, r^* + 1, b^*)$. Observe that a^* is an exit of (T, B') by the condition of the if-block and b^* is an exit of (T, B^*) by the definition of a destination. It remains to check that $\mathsf{dist}_T(a^*, B') = 1$ and $\mathsf{dist}_T(a^*, B^*) = r^* + 1$. The former is immediate since a^* is adjacent to B' in T. The latter is true since $r^* = \mathsf{dist}_T(a^*, b^*)$ and B^* is adjacent to b^* in T where the (unique) path from a^* to B^* goes via b^*. ∎

Lemma 9.18. $\textsc{ExitWalk}(T, a^*, B, \ell, a_0)$ *(Algorithm 9.2) is well-defined, terminates, and returns a closed walk in H of length at least 3 from a_0 to itself.*

Proof. First consider the case $k = 0$. Clearly, Algorithm 9.2 terminates and is well-defined. It returns $\mathrm{W}_C(a_0)$, which is a cycle from a_0 to itself of length at least 3. Now consider the case where $k \geq 1$. Note that with each recursive call of $\textsc{ExitWalk}(\cdot)$ the value of the parameter ℓ increases. Since ℓ corresponds to the distance between a^* and B in the finite graph T, Algorithm 9.2 terminates.

We now show that Algorithm 9.2 returns a closed walk of length at least 3 from a_0 to itself. If, for $i \in [k]$, $\textsc{ExitWalk}(T, a^*, B_i, \ell + r_i + 1, b_i)$ returns a closed walk from b_i to itself of length at least 3 then $\mathrm{P}_H(a_i, b_i) + \textsc{ExitWalk}(T, a^*, B_i, \ell + r_i + 1, b_i) + \mathrm{P}_H(b_i, a_i) + \mathrm{W}_C(a_i, a_{i+1 \bmod k+1})$ is a walk from a_i to $a_{i+1 \bmod k+1}$ of length at least 3. Thus, $\textsc{ExitWalk}(T, a^*, B, \ell, a_0)$ returns a closed walk from a_0 to itself of length at least 3. Since Algorithm 9.2 terminates, it reaches the base of the recursion, i.e., the case $k = 0$, at some point, and we have already verified that the base case returns a closed walk of length at least 3, as required.

Finally, we show that Algorithm 9.2 is well-defined. By Lemma 9.17 all subroutine calls have feasible inputs. Also observe that all concatenation operations are well-

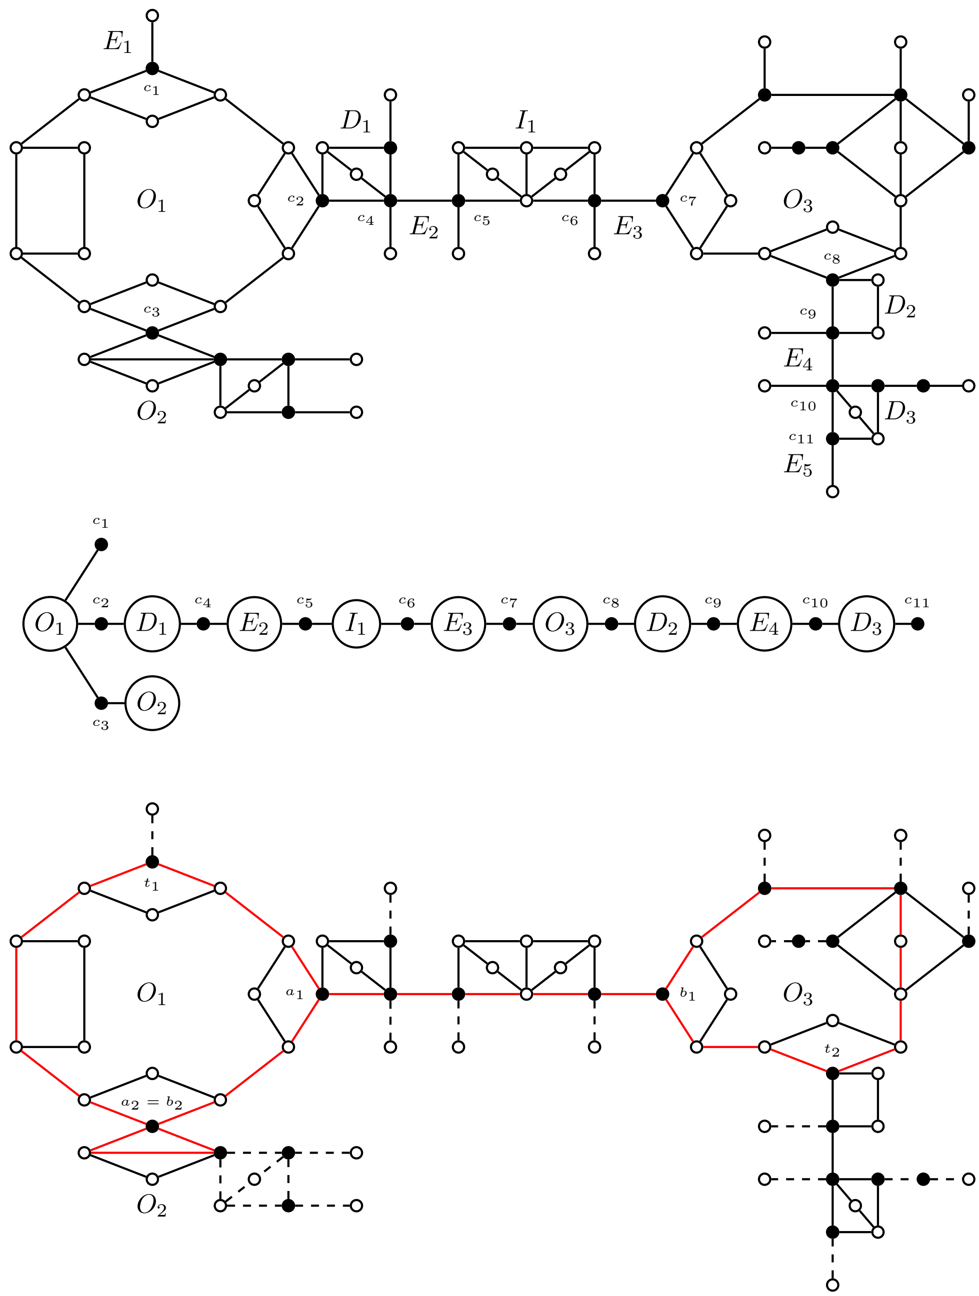

Fig. 10: *A graph H, a closed suitable subtree T of* $\mathsf{BC}(H)$ *with a block O_1 that is an obstruction, and* $\textsc{Walk}(T, O_1)$*.*
(Top) An involution-free and K_4-minor-free graph H such that every biconnected component is an edge, a diamond, an impasse or an obstruction. Articulation points are depicted as filled vertices.
(Center) A *closed* and *suitable* subtree T of the block-cut tree of H, rooted at O_1. Note that every cut vertex of T that is a leaf (i.e., c_1 or c_{11}) is O_1-closed in $\mathsf{BC}(H)$.
(Bottom) Solid lines are contained in the subgraph of H induced by $V(T)$, while dashed lines are not. The red closed walk is the output of $\textsc{Walk}(T, O_1)$. Observe that a_1 and a_2 are exits of (T, O_1) with destinations (b_1, O_3) and (b_2, O_2), respectively, and that t_1 and t_2 are attachment points of (T, O_1) and (T, O_3), respectively.

defined since, for each $i \in [k]$, $\textsc{ExitWalk}(T, a^*, B_i, \ell + r_i + 1, b_i)$ returns a closed walk from b_i to itself. $\square$

LEMMA 9.19. $\textsc{Walk}(T, B')$ *(Algorithm 9.3) terminates, is well-defined, and returns a closed walk in H of length q, where $q \geq 3$ and $q \neq 4$.*

Proof. Since Algorithm 9.2 terminates (Lemma 9.18), it is immediate that Algorithm 9.3 terminates. If there is no exit of (T, B') then $\textsc{Walk}(T, B')$ returns $C(B', \Gamma_T(B'))$ which is a cycle in $\mathrm{Cy}(B')$ by Definition 9.1 and hence has length at least 3, but not 4, by Definition 8.6. If there is an exit a^* of (T, B') then, by Lemma 9.17 all subroutine calls have feasible inputs. By Lemma 9.18, the algorithm $\textsc{ExitWalk}(T, a^*, B', 1, a^*)$ returns a closed walk from a^* to itself, of length at least 3, and $\textsc{ExitWalk}(T, a^*, B^*, r^* + 1, b^*)$ returns a closed walk from b^* to itself, also of length at least 3. It follows that the concatenations in the if-block are well-defined and therefore that Algorithm 9.3 is well-defined. Also, $\textsc{ExitWalk}(T, a^*, B', 1, a^*) + \mathrm{P}_H(a^*, b^*) + \textsc{ExitWalk}(T, a^*, B^*, r^* + 1, b^*) + \mathrm{P}_H(b^*, a^*)$ is a closed walk from a^* to itself of length $q \geq 6$. $\square$

OBSERVATION 9.20. $\textsc{Walk}(T, B')$ *(Algorithm 9.3) outputs a closed walk W. If (T, B') has no exit then $W = C$ for a cycle $C \in Cy(B')$. Otherwise, the following holds. For a positive integer j, there are obstructions $B'_0, \ldots, B'_j$ such that W is of the form $W = Q_0 + P_0 + Q_1 + P_1 \cdots + Q_j + P_j$ where Q_i and P_i satisfy the following properties for all $i \in \{0, \ldots, j\}$.*

- *Let $C_i = C(B'_i, \Gamma_T(B'_i))$. Then there are vertices a and a' in B'_i such that Q_i is of the form $\mathrm{W}_{C_i}(a)$ or $\mathrm{W}_{C_i}(a, a')$. Either way, all vertices of Q_i are in B'_i. Furthermore, only the endpoints of Q_i are exits of (T, B'_i).*
- *There is an exit a of (T, B'_i) with a destination $(b, B'_{i+1 \bmod (j+1)})$ such that P_i is a path of the form $\mathrm{P}_H(a, b)$. Hence, by Definition 9.14 and Lemma 9.13, the endpoints of P_i are the only vertices of P_i that are part of an obstruction.*
- *The obstruction B'_i is distinct from the obstruction $B'_{i+1 \bmod (j+1)}$.*

Explanation of Observation 9.20. If (T, B') has no exit then the result follows directly from the definition of $C(B, A)$ for a suitable connector (B, A) of an obstruction B (Definition 9.1).

For the remaining case, we will prove that $\textsc{ExitWalk}(T, a^*, B, \ell, a_0)$ (Algorithm 9.2) outputs a closed walk W'. For a positive integer q, there are obstructions $B''_0, \ldots, B''_q$ such that W' is of the form $W' = Q'_0 + \sum_{i=1}^{q} (P'_i + Q'_i)$ where Q'_i and P'_i satisfy the following properties for all $i \in \{0, \ldots, q\}$.

- Let $C_i = C(B''_i, \Gamma_T(B''_i))$. Then there are vertices a and a' in B''_i such that Q'_i is of the form $\mathrm{W}_{C_i}(a)$ or $\mathrm{W}_{C_i}(a, a')$. Either way, all vertices of Q'_i are in B''_i. Furthermore, only the endpoints of Q'_i are exits of (T, B''_i).
- There is an exit a of (T, B''_i) with a destination $(b, B''_{i+1 \bmod (j+1)})$ such that P'_i is a path of the form $\mathrm{P}_H(a, b)$. Hence, by Definition 9.14, the endpoints of P'_i are the only vertices of P'_i that are part of an obstruction.
- The obstruction B''_i is distinct from the obstruction $B''_{i+1 \bmod (q+1)}$.

The proof is by induction on the recursion depth. If $\textsc{ExitWalk}(T, a^*, B, \ell, a_0)$ makes no recursive calls then the variable "k" is equal to 0 and we also set $q = 0$. In this case, B''_0 is equal to B, and the first property follows easily (the others are vacuous). Otherwise, k is positive and (T, B) has exits $\{a_0, \ldots, a_k\}$ where $a_1, \ldots, a_k$ have destinations $(b_1, B_1), \ldots, (b_k, B_k)$. Note that $B_1, \ldots, B_k$ are disjoint from B and from each other. Once again, B''_0 is B. Q'_0 is $W_C(a_0, a_1)$, as defined in the algorithm. Then B''_1 is B_1, and P'_1 is $P_H(a_1, b_1)$ as defined in the algorithm. The rest follows by induction, and examination of the algorithm, using the fact that the block-cut tree is a tree.

Given this fact for Algorithm 9.2, we obtain the conclusion for Algorithm 9.3 by putting together the pieces in the output W. This completes our explanation of Observation 9.20.

LEMMA 9.21. *Let H be a connected graph. Let T be a closed suitable subtree of $\mathsf{BC}(H)$. Let B' be an obstruction that is a block of T. Let $W = (w_0, , \dots, w_{q-1}, w_0)$ be the output of* WALK(T, B') *(Algorithm 9.3). Then, for each $i \in \{0, \dots, q-1\}$, w_i and $w_{i+2 \mod q}$ are distinct.*

Proof. All indices in this proof are considered to be modulo q. For any $i \in \{0, \dots, q-1\}$, our goal is to show $w_i \neq w_{i+2}$. We make a case distinction based on Observation 9.20.

- If W is a cycle $C \in \mathrm{Cy}(B')$ then $w_i \neq w_{i+2}$ is immediate.
- Otherwise, for a positive integer j, W is of the form $W = Q_0 + P_0 + Q_1 + P_1 \cdots + Q_j + P_j$ with the properties stated in Observation 9.20. We consider the walk (w_i, w_{i+1}, w_{i+2}).
  - If for some $\ell \in [j]$, (w_i, w_{i+1}, w_{i+2}) is a subwalk of Q_ℓ then $w_i \neq w_{i+2}$ since, by Observation 9.20, Q_i is a subwalk of a cycle.
  - If for some $\ell \in [j]$, (w_i, w_{i+1}, w_{i+2}) is a subwalk of P_ℓ then, since P_ℓ is a path, we have $w_i \neq w_{i+2}$.
  - Otherwise, by Observation 9.20, there is no biconnected component that contains both w_i and w_{i+2} and consequently $w_i \neq w_{i+2}$. □

The following lemma establishes (a stronger version of) the properties (L5.5.2) and (L5.5.3) for the walk returned by Algorithm 9.3, as required by Lemma 5.6.

LEMMA 9.22. *Let H be a connected K_4-minor-free graph. Let T be a closed suitable subtree of $\mathsf{BC}(H)$. Let B' be an obstruction that is a block of T. Let $W = (w_0, \dots, w_{q-1}, w_0)$ be the output of* WALK(T, B') *(Algorithm 9.3). By Lemma 9.19, W is a closed walk and $q \geq 3$. For each $i \in \{0, \dots, q-1\}$, let $W_i = N_{W,H}(w_i)$ (Definition 4.2). If H has no hardness gadget then the following statement holds:*

$$\text{If } u \in W_{i-1 \bmod q} \text{ and } v \in W_{i+1 \bmod q} \text{ then } \Gamma_H(u) \cap \Gamma_H(v) = W_i.$$

Proof. All indices in this proof are considered to be modulo q. Let $i \in \{0, \dots, q-1\}$, $u \in W_{i-1}$ and $v \in W_{i+1}$. Our goal is to show that $\Gamma_H(u) \cap \Gamma_H(v) = W_i$. We split the proof into two cases (Claims A and B).

Claim A: *If there is no biconnected component of H that contains both w_{i-1} and w_{i+1} then $\Gamma_H(u) \cap \Gamma_H(v) = W_i$.*

Proof: If there is no biconnected component that contains both w_{i-1} and w_{i+1} then, by the definition of W_i, $W_i = \{w_i\}$. Since w_{i-1} and w_{i+1} are not in the same biconnected component every path from w_{i-1} to w_{i+1} goes through w_i. There is a path from w_{i-1} to u via w_{i-2} and there is a path from v to w_{i+1} via w_{i+2}. Since w_{i-2} and w_{i+2} are distinct from w_i by Lemma 9.21 these paths do not go through w_i. Hence every path from u to v also goes through w_i. Thus, there is no biconnected component that contains both u and v. Hence, $\Gamma_H(u) \cap \Gamma_H(v) = \{w_i\} = W_i$, as required. This concludes the proof of Claim A. ■

Claim B: *If there is a biconnected component B such that w_{i-1} and w_{i+1} are in B then $\Gamma_H(u) \cap \Gamma_H(v) = W_i$.*

Proof: By Lemma 9.21, $w_{i-1} \neq w_{i+1}$. This together with the fact that w_i is adjacent to both w_{i-1} and w_{i+1} implies that w_i is also in B. If $u = w_{i-1}$ then it is trivial that u is in B. If $u \neq w_{i-1}$ then $|W_{i-1}| > 1$. By the fact that $W_{i-1} = \Gamma_H(w_{i-2}) \cap \Gamma_H(w_i)$ and the fact that both w_{i-1} and w_i are in B, it follows that $W_{i-1} \subseteq V(B)$ and that w_{i-2} is in B. Thus, we have established that u is in B. Analogously, v is in B. We state this formally so we can refer to it.

Fact 1: *If $|W_{i-1}| > 1$ then every vertex in $W_{i-1} \cup \{w_{i-2}\}$ is in B. Similarly, if $|W_{i+1}| > 1$ then every vertex in $W_{i+1} \cup \{w_{i+2}\}$ is in B. Consequently, both u and v are in B.*

If $u = w_{i-1}$ and $v = w_{i+1}$ then $\Gamma_H(u) \cap \Gamma_H(v) = \Gamma_H(w_{i-1}) \cap \Gamma_H(w_{i+1}) = W_i$, as required.

Therefore, we assume for the rest of the proof that $u \neq w_{i-1}$ (the case $v \neq w_{i+1}$ is symmetric). By Fact 1, the walk $(w_{i-2}, w_{i-1}, w_i, w_{i+1})$ is in B. We show that B is an obstruction. Suppose, for contradiction, that B is an edge, diamond or impasse. Then by Observation 9.20, there are cut-vertices a and b such that the walk $(w_{i-2}, w_{i-1}, w_i, w_{i+1})$ is a subpath of $\mathrm{P}_H(a, b)$. This contradicts Lemma 9.13, which states that no four consecutive vertices of this path are part of the same biconnected component.

Thus, we have established that $(w_{i-2}, w_{i-1}, w_i, w_{i+1})$ is a walk in the obstruction B. By Observation 9.20 and the definition of $\mathrm{W}_C(\cdot)$ (Definition 9.16), it is a subwalk of some cycle $C \in \mathrm{Cy}(B)$ following the order $D(C)$. It follows that $W_{i-1} = N_{C,H}(w_{i-1})$ and $W_i = N_{C,H}(w_i)$. By Corollary 8.5, from the fact that H has no hardness gadget and $|W_{i-1}| > 1$ it follows that $W_i = \{w_i\}$. Let ℓ be the length of C. Since $C \in \mathrm{Cy}(B)$ we have $\ell = 3$ or $\ell > 4$. We make a case distinction depending on ℓ.

- Suppose $\ell = 3$ so $w_{i-2} = w_{i+1}$. Suppose, for contradiction that $|W_{i+1}| > 1$. Then by Fact 1, w_{i+2} is also in B and $(w_{i-2}, w_{i-1}, w_i, w_{i+1}, w_{i+2})$ is a subwalk of $\mathrm{W}_C(\cdot)$. This gives a contradiction to the fact that all vertices of $\mathrm{W}_C(\cdot)$, apart from possibly its endpoints, are distinct (see Definition 9.16). Therefore, $W_{i+1} = \{w_{i+1}\}$ and consequently $v = w_{i+1}$. Since $u \neq w_{i-1}$, (v, u, w_i, v) and (v, w_{i-1}, w_i, v) are two distinct triangles that share the edge $\{w_i, v\}$. By Lemma 6.9, since H has no hardness gadget, u and v have no common neighbour other than w_i.
- Suppose $\ell > 4$. Apply Lemma 8.3 to the cycle C and the index $i - 1$. This shows that there is a separation (A_1, A_2) of H such that $C \setminus \{w_{i-1}\} \subseteq A_1$, $W_{i-1} = N_{C,H}(w_{i-1}) \subseteq A_2$, and $A_1 \cap A_2 = \{w_{i-2}, w_i\}$. Since $u \in A_2$, u is not adjacent to any vertex in $C \setminus \{w_{i-2}, w_{i-1}, w_i\}$. By the definition of $\mathrm{Cy}(B)$ (Definition 8.6) C is an induced cycle of B, so the cycle C' that is obtained from C by replacing w_{i-1} with u is also an induced cycle of B. Also, C' has length $\ell > 4$. Since H has no hardness gadget, we can apply Corollary 8.5 to obtain (using the fact that $N_{C',H}(u) = N_{C,H}(w_{i-1}) = W_{i-1}$ has cardinality greater than 1) that $|N_{C',H}(w_i)| = 1$. By definition, $N_{C',H}(w_i) = \Gamma_H(u) \cap \Gamma_H(w_{i+1})$, so $\Gamma_H(u) \cap \Gamma_H(w_{i+1}) = \{w_i\}$.
  - If $|W_{i+1}| = 1$ then $v = w_{i+1}$, so we are finished.
  - Suppose that $|W_{i+1}| > 1$. By Fact 1, w_{i+1} and w_{i+2} are in B, and consequently the walk $(w_{i-2}, u, w_i, w_{i+1}, w_{i+2})$ is a subwalk of C'. It follows that $W_{i+1} = N_{C',H}(w_{i+1})$. Apply Lemma 8.3 to the cycle C' and the index $i + 1$. This shows that there is a separation (A_3, A_4) of H such that $C' \setminus \{w_{i+1}\} \subseteq A_3$, $W_{i+1} \subseteq A_4$, and $A_3 \cap A_4 = \{w_i, w_{i+2}\}$.

Since $v \in A_4$, v is not adjacent to any vertex in $C' \setminus \{w_i, w_{i+1}, w_{i+2}\}$. So the cycle C'' that is obtained from C' by replacing w_{i+1} with v is also an induced cycle of B with length $\ell > 4$. Since H has no hardness gadget, we can apply Corollary 8.5 to obtain (using the fact that $N_{C'',H}(v) = W_{i+1}$ has cardinality greater than 1) that $|N_{C'',H}(w_i)| = 1$. By definition, $N_{C'',H}(w_i) = \Gamma_H(u) \cap \Gamma_H(v)$, so we are finished. □

This concludes the proof of Claim B. ■

The lemma follows immediately from Claim A and Claim B.

The following lemma establishes property (L5.5.4) for the walk returned by Algorithm 9.3, as required by Lemma 5.6.

LEMMA 9.23. *Let H be a connected graph. Let T be a closed suitable subtree of $\mathsf{BC}(H)$. Let B' be an obstruction that is a block of T. Let $W = (w_0, , \ldots, w_{q-1}, w_0)$ be the output of* WALK(T, B') *(Algorithm 9.3). By Lemma 9.19, W is a closed walk and $q \geq 3$. For each $i \in \{0, \ldots, q-1\}$, let $W_i = N_{W,H}(w_i)$. Then there exists no closed walk $D = (d_0, \ldots, d_{q-1}, d_0)$ with $d_i \in \Gamma_H(W_i) \setminus (W_{i-1} \cup W_{i+1})$ for all i (indices taken modulo q).*

Proof. Assume for the sake of contradiction that such a walk D exists. We distinguish two cases:

(I) W is not entirely contained in a single biconnected component of H. In this case, there is an index i such that no biconnected component contains both w_{i-1} and w_{i+1}. Now consider $d_{i-1} \in \Gamma_H(w_{i-1})$ and $d_{i+1} \in \Gamma_H(w_{i+1})$. Note that $d_{i-1} \neq w_i$ and $d_{i+1} \neq w_i$ by the specification of D. Note further that $d_i \neq w_i$ as we do not allow self-loops in H. Consequently, there are two internally vertex disjoint 2-paths from w_{i-1} to w_{i+1}; one passes through d_{i-1}, d_i and d_{i+1}; and the other passes through w_i. This is a contradiction to the fact that no biconnected component contains both w_{i-1} and w_{i+1}.

(II) W is entirely contained in a biconnected component B. By Observation 9.20, the only possibility for this to be true is that B is an obstruction and W is a cycle in $\mathrm{Cy}(B)$. By the definition of obstructions (and $\mathrm{Cy}(B)$), W is thus an induced cycle of length q such that $q \geq 3$ and $q \neq 4$. The lemma will follow easily from the following claim.

Claim A: $D \cap (\bigcup_i W_i) = \emptyset$.

Proof: Assume for contradiction that $d_i \in W_j$ for some indices i and j. Note that $j \notin \{i-1, i+1\}$ by the specification of D. We cannot have $j = i$ since $d_i \in \Gamma_H(W_i)$ so $d_i \notin W_i$ (otherwise H has a self-loop). Since $j \notin \{i-1, i, i+1\}$, we have $q \geq 5$. We will show that H has a K_4-minor. Since d_i is adjacent to w_i, and it is not equal to w_{i-1} or w_{i+1} and since W is an induced cycle in B, we conclude that d_i is distinct from the vertices of W. H therefore contains a K_4-minor containing the vertex d_i and its three neighbours w_i, w_{j-1} and w_{j+1}. There are disjoint paths between these three vertices along the cycle W and d_i is not on these paths. See the following illustration.

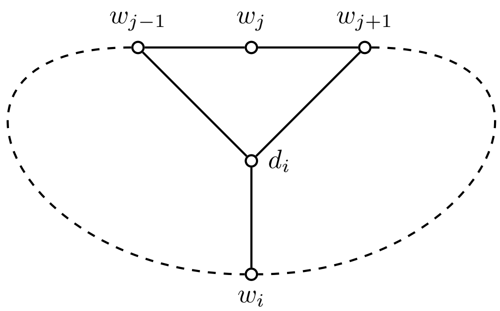

This concludes the proof of Claim A. ■

We have assumed for contradiction that D exists, and proved Claim A. We obtain the contradiction by using Claim A to construct a K_4-minor in H. Claim A demonstrates that $W \cap D = \emptyset$. Now contract the walk D to a single vertex. This yields a vertex $d \notin W$ which is adjacent to all vertices of W. As W has length at least 3, we have found a K_4-minor as promised. □

LEMMA 9.24. *Let H be a connected K_4-minor-free graph. Let T be a closed suitable subtree of $\mathsf{BC}(H)$. Let B' be an obstruction that is a block of T. Then H has a hardness gadget.*

Proof. Let $W = (w_0, \ldots, w_{q-1}, w_0)$ be the output of WALK(T, B'). According to Lemma 9.19, W is a closed walk with $q \geq 3$ and $q \neq 4$. Our goal is to use Lemma 5.6 to show that H has a hardness gadget. To this end, we identify the sets $\mathcal{C}_i$ of Lemma 5.6 with the sets $W_i = N_{W,H}(w_i)$. Let S be the set of all i such that W_i has even cardinality.

Claim A: *For every $i \in S$, there is an obstruction O_i such that the following hold for $C_i = C(O_i, \Gamma_T(O_i))$.*

- *$w_{i-1}, w_i, w_{i+1} \in C_i$.*
- *$W_i = N_{C_i,H}(w_i)$.*
- *Every vertex in W_i has degree 2 in O_i.*
- *w_i is an attachment point of (T, O_i).*

Proof: Fix $i \in S$. By the definition of W_i and the fact that $|W_i| > 1$, there is a biconnected component O_i of H that contains w_{i-1}, W_i, and w_{i+1}. Suppose, for contradiction, that O_i is an edge, diamond or impasse, then, by Observation 9.20, the walk (w_{i-1}, w_i, w_{i+1}) is a subpath of a path of the form $\mathrm{P}_H(\cdot)$. However, since $|W_i| \geq 2$, w_{i-1} and w_{i+1} have at least 2 common neighbours in H, this is a contradiction to Lemma 9.13.

We have established that O_i is an obstruction. Since T is a suitable subtree, $(O_i, \Gamma_T(O_i))$ is a suitable connector and, by Definition 9.1, $C_i = C(O_i, \Gamma_T(O_i))$ is a cycle with $\Gamma_T(O_i) = \{c \in C_i \mid \text{the cardinality of } N_{C_i,H}(c) \text{ is even}\}$. By Observation 9.20, (w_{i-1}, w_i, w_{i+1}) is a subwalk of a walk of the form $\mathrm{W}_{C_i}(\cdot)$. It follows that $w_{i-1}, w_i, w_{i+1} \in C_i$ and $W_i = N_{C_i,H}(w_i)$, as required.

As the cardinality of W_i is even and $C_i \in \mathrm{Cy}(O_i)$, by Definition 8.6, every vertex in W_i has degree 2 in O_i, as required.

The fact that the cardinality of W_i is even also implies that $w_i \in \Gamma_T(O_i)$ (since $\Gamma_T(O_i) = \{c \in C_i \mid \text{the cardinality of } N_{C_i,H}(c) \text{ is even}\}$). Thus, by Definition 9.14, w_i is either an exit or an attachment point of (T, O_i). However, by Observation 9.20, only the endpoints of $\mathrm{W}_{C_i}(\cdot)$ are exits, which means that w_i is an attachment point, as required. This finishes the proof of Claim A. ■

In the remainder of this proof, for each $i \in S$, let O_i and C_i be as stated in Claim A. Next we use the fact that, for each $i \in S$, w_i is an attachment point of (T, O_i) to define a gadget $(\hat{J}_i, \hat{z}_i)$. Those gadgets will be used in the construction of the gadgets $(J_0, z_0), \ldots, (J_{q-1}, z_{q-1})$ required by Lemma 5.6. Recall that, by definition of attachment points (Definition 9.14), for each $i \in S$, there is a (unique) maximal-length proper obstruction-free path $P_i = (b_0^i, B_1^i, b_1^i, B_2^i, \ldots, B_{q_i}^i, b_{q_i}^i)$ in T such that $w_i = b_{q_i}^i$ and b_0^i is a leaf in T. As T is closed, we obtain that, for some block R of T, the vertex b_0^i is R-closed, i.e., b_0^i has precisely one descendant in $\mathsf{BC}(H)$ with respect to $<_R$. Moreover, this descendant must be an edge. We distinguish whether $q_i = 0$ or $q_i \geq 1$:

$q_i = 0$: We have $w_i = b_{q_i}^i = b_0^i$. Since $b_0^i = w_i$ is R-closed, Lemma 9.7 ensures that it is also O_i-closed. Consequently, w_i has precisely three neighbours in H: The two neighbours in O_i (which are w_{i-1} and w_{i+1} — these are distinct by Lemma 9.21), as well as the other endpoint ℓ_i of the edge that is the unique descendant of w_i in $\mathsf{BC}(H)$.

We define $\hat{J}_i$ to be a single edge, one endpoint of which is z_i, and the other endpoint of which is pinned to w_i. Observe that

$$\{v \in V(H) \mid \left|\hom\Big((\hat{J}_i, \hat{z}_i) \to (H, v)\Big)\right| \text{ is odd.}\} = \{w_{i-1}, w_{i+1}, \ell_i\}\,.$$

This concludes the definition of $(\hat{J}_i, \hat{z}_i)$ in the case that $q_i = 0$.

$q_i \geq 1$: By Lemma 9.7, b_0^i is B_0^i-closed. It follows that $|\Gamma_H(b_0^i) \setminus B_0^i| = 1$. Since T is a suitable subtree and P_i is obstruction-free, for each $j \in [q_i]$, $(B_j^i, \{b_{j-1}^i, b_j^i\})$ is a suitable connector in H and B_j^i is an edge, diamond or impasse. Thus, we can invoke Lemma 7.14. We obtain that at least one of the following is true:

- H has a hardness gadget.
- $B_{q_i}^i$ is an edge or a diamond and $(L_i, b_{q_i}^i)$ is a good start in $B_{q_i}^i$ but not a good stop in $B_{q_i}^i$, where $L_i = \{b_{q_i-1}^i\}$.
- $B_{q_i}^i$ is an impasse, and $(L_i, b_{q_i}^i)$ is a good start in $B_{q_i}^i$ but not a good stop in $B_{q_i}^i$, where $L_i = \{d_i\}$ and d_i is the unique common neighbour of $b_{q_i-1}^i$ and $b_{q_i}^i$ in H.

We are done in the first case, so suppose that one of other cases applies. By definition of good starts, we thus obtain a gadget $(\hat{J}_i, \hat{z}_i)$ such that, for $R_i = \Gamma_H(b_{q_i}^i) \setminus V(B_{q_i}^i)$,

$$\{v \in V(H) \mid \left|\hom\Big((\hat{J}_i, \hat{z}_i) \to (H, v)\Big)\right| \text{ is odd.}\} = L_i \cup R_i.$$

Note that L_i and R_i are disjoint. Further, recall that $w_i = b_{q_i}^i$ and therefore $R_i \cap V(O_i) = \{w_{i-1}, w_{i+1}\}$ (by Claim A). As $(L_i, b_{q_i}^i)$ is not a good stop in $B_{q_i}^i$, we have that R_i is of even cardinality, and thus $L_i \cup R_i$ is of odd cardinality. This concludes the definition of $(\hat{J}_i, \hat{z}_i)$ in the case that $q_i \geq 1$.

We now state the previously-established crucial property of the gadgets $(\hat{J}_i, \hat{z}_i)$ (which applies for all q_i, unless H has a hardness gadget).

Fact 1: *For every $i \in S$, there is a gadget $(\hat{J}_i, \hat{z}_i)$ such that the set*

$$\hat{\Omega}_i = \{v \in V(H) \mid \left|\hom\Big((\hat{J}_i, \hat{z}_i) \to (H, v)\Big)\right| \textit{ is odd.}\}$$

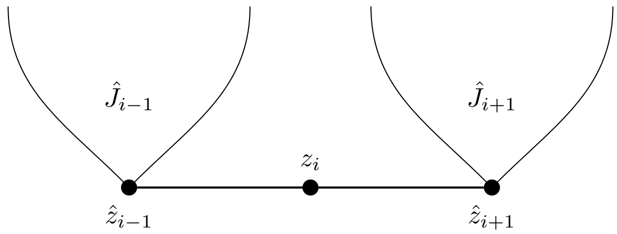

Fig. 11: The gadget (J_i, z_i) for the case $i-1 \in S$ and $i+1 \in S$

is a subset of $\Gamma_H(w_i)$*, has odd cardinality, and contains precisely two vertices of* O_i *— the vertices* w_{i-1} *and* w_{i+1}*.*

We proceed by defining for each $i \in \{0, \ldots, q-1\}$ the gadgets (J_i, z_i) needed for Lemma 5.6. The definition of (J_i, z_i) depends on whether or not $i-1$ and $i+1$ are in S. (J_i, z_i) always contains the vertex z_i. Additionally, for $j \in \{i-1, i+1\}$, if $j \in S$ then (J_i, z_i) also contains a copy of the gadget $(\hat{J}_j, \hat{z}_j)$ and z_i is adjacent to $\hat{z}_j$. We provide an illustration of (J_i, z_i) for the case $i-1 \in S$ and $i+1 \in S$ in Figure 11.

Following the notation of Lemma 5.6, we set (for every $i \in \{0, \ldots, q-1\}$):

$$\Omega_i = \{v \in V(H) \mid \ \ |\mathrm{hom}((J_i, z_i) \to (H, v))| \text{ is odd.}\}$$

Claim B: *For all* $i \in \{0, \ldots, q-1\}$ *the following holds true, unless* H *has a hardness gadget; indices are taken modulo* q*:*

- *If* $i-1 \in S$ *then* $W_{i-1} \cap \Omega_i = \{w_{i-1}\}$*, and if* $i-1 \notin S$ *then* $W_{i-1} \cap \Omega_i = W_{i-1}$*.*
- *If* $i+1 \in S$ *then* $W_{i+1} \cap \Omega_i = \{w_{i+1}\}$*, and if* $i+1 \notin S$ *then* $W_{i+1} \cap \Omega_i = W_{i+1}$*.*

So, $W_{i-1} \cap \Omega_i$ *and* $W_{i+1} \cap \Omega_i$ *have odd cardinality.*

Proof: We only show the first item; the second one is symmetric. We distinguish whether or not $i-1 \in S$:

(I) If $i-1 \in S$ then, by Claim A, w_{i-1} and w_i are contained in the cycle C_{i-1}, and $W_{i-1} = N_{C_{i-1},H}(w_{i-1})$. Since $C_{i-1} \in \mathrm{Cy}(O_{i-1})$, C_{i-1} is an induced cycle in O_{i-1} (hence in H) and is not a square. Therefore we can apply Corollary 8.5. It follows that H either has a hardness gadget, in which case we are done, or $\left|N_{C_{i-1},H}(w_i)\right| = 1$, i.e., $N_{C_{i-1},H}(w_i) = \{w_i\}$. This implies that w_i is not an exit of (T, O_{i-1}), and thus $w_{i+1} \in C_{i-1}$ and consequently $W_i = N_{C_{i-1},H}(w_i) = \{w_i\}$.
We are now able to prove that $W_{i-1} \cap \Omega_i = \{w_{i-1}\}$. First, for $v \in W_{i-1} \setminus \{w_{i-1}\}$, we show that $v \notin \Omega_i$ and hence $v \notin W_{i-1} \cap \Omega_i$. By the construction of J_i, it suffices to show that there is an even number of vertices in $\hat{\Omega}_{i-1}$ that are adjacent to v. Recall from Fact 1 that $\hat{\Omega}_{i-1} \subseteq \Gamma_H(w_{i-1})$. The vertex v has precisely two common neighbours with w_{i-1}, namely w_{i-2} and w_i (any others would lead to a K_4-minor in H induced by the vertices $\{v, w_{i-2}, w_{i-1}, w_i\}$). By Fact 1, we know that both of these are in $\hat{\Omega}_{i-1}$ and hence that there are two vertices in $\hat{\Omega}_{i-1}$ that are adjacent to v, as required.
It remains to show that $w_{i-1} \in \Omega_i$ and hence $w_{i-1} \in W_{i-1} \cap \Omega_i$. By the construction of J_i, it suffices to show that there is an odd number of vertices in $\hat{\Omega}_{i-1}$ that are adjacent to w_{i-1}, and, in case $i+1 \in S$, that there is an odd number of vertices in $\hat{\Omega}_{i+1}$ that are adjacent to w_{i-1}.

- By Fact 1, $\hat{\Omega}_{i-1} \subseteq \Gamma_H(w_{i-1})$. Hence every element of $\hat{\Omega}_{i-1}$ is adjacent to w_{i-1}. By Fact 1, $\hat{\Omega}_{i-1}$ has odd cardinality, as required.
- Suppose that $i+1 \in S$. By Fact 1, we have $\hat{\Omega}_{i+1} \subseteq \Gamma_H(w_{i+1})$ and $w_i \in \hat{\Omega}_{i+1}$. Furthermore, w_i is the only common neighbour of w_{i-1} and w_{i+1} in H by the fact that $W_i = \{w_i\}$. Hence w_i is the only vertex in $\hat{\Omega}_{i+1}$ that is adjacent to w_{i-1}, as required.

(II) Consider $i-1 \notin S$. Our goal is to show that $W_{i-1} \cap \Omega_i = W_{i-1}$. If $i+1 \notin S$, then J_i contains only the single vertex z_i and $\Omega_i = V(H)$ and we are finished. Hence we can assume $i+1 \in S$. We first proceed as in Case (I) to obtain either a hardness gadget or $W_i = \{w_i\}$. By Claim A, w_i, w_{i+1} and w_{i+2} are contained in the cycle C_{i+1}, and $W_{i+1} = N_{C_{i+1},H}(w_{i+1})$. Since $C_{i+1} \in \mathrm{Cy}(O_{i+1})$, C_{i+1} is induced and not a square and therefore we can apply Corollary 8.5. It follows that H either has a hardness gadget, in which case we are done, or $\left|N_{C_{i+1},H}(w_i)\right| = 1$, i.e., $N_{C_{i+1},H}(w_i) = \{w_i\}$. This implies that w_i is not an exit of (T, O_{i+1}), and thus $w_{i-1} \in C_{i+1}$ and consequently $W_i = N_{C_{i+1},H}(w_i) = \{w_i\}$.
In order to show that $W_{i-1} \cap \Omega_i = W_{i-1}$ we show, for each $v \in W_{i-1}$, that $v \in \Omega_i$. By the construction of J_i ($i-1 \notin S$, $i+1 \in S$), it suffices to show that there is an odd number of vertices in $\hat{\Omega}_{i+1}$ that are adjacent to v. Recall from Fact 1 that $\hat{\Omega}_{i+1} \subseteq \Gamma_H(w_{i+1})$. By the fact that $v \in W_{i-1}$ and $w_{i+1} \in W_{i+1}$, from Lemma 9.22 we obtain that $\Gamma_H(v) \cap \Gamma_H(w_{i+1}) = W_i$. We have already established that $W_i = \{w_i\}$ and hence w_i is the only vertex in $\hat{\Omega}_{i+1}$ that is adjacent to w_{i-1}, as required.

This concludes the proof of Case (II) and of Claim B. ■

We prove one final claim before we can apply Lemma 5.6:

Claim C: *Unless H has a hardness gadget, there exists $k \in \{0, \ldots, q-1\}$ such that both of the following are true; indices are taken modulo q:*
- *There are no edges between W_k and W_{k+3}.*
- *$(W_k \cup W_{k+2}) \cap \Omega_{k+1}$ and $(W_{k+1} \cup W_{k+3}) \cap \Omega_{k+2}$ are of even cardinality.*

Proof: We distinguish two cases.

(I) There is a biconnected component B that contains W. Consequently, by Observation 9.20, there is a cycle $C \in \mathrm{Cy}(B)$ such that $W = C$. Since $C \in \mathrm{Cy}(B)$ it has length $q = 3$ or $q \geq 5$. In this case, we choose $k = 0$. We first show that there is no edge between W_k and W_{k+3}:
- If $q = 3$, we show that for $u, v \in W_0$ there cannot be an edge between u and v. If $u = v$ there cannot be an edge since we do not allow self-loops in H. If $u \neq v$ there cannot be an edge, as otherwise u, v, w_1, w_2 induce a K_4-minor in H, contradicting the fact that H has none. Thus, there are no edges between W_0 and $W_{3 \bmod q} = W_0$.
- If $q \geq 5$, consider W_0 and $W_3 = W_{3 \bmod q}$. If $|W_0| = |W_3| = 1$ then there are no edges between W_0 and W_3 since C is induced by the definition of obstruction (Definition 8.6). So suppose $|W_0| > 1$ (the case $|W_3| > 1$ is symmetric). Since C is an induced cycle of length $q > 4$ in a biconnected graph B, we can apply Lemma 8.3 to find a separation (A, A') of H such that $C \setminus \{w_0\} \subseteq A$, $W_0 \subseteq A'$and $A \cap A' = \{w_q, w_1\}$. Since all of the vertices in $\bigcup_{i=1}^{q-1} W_i$ have neighbours in $C \setminus \{w_0\}$, this implies that w_{q-1} and w_1 are the only vertices in $\bigcup_{i=1}^{q-1} W_i$ that are adjacent to vertices in W_0. However, by Lemma 8.2, $W_0, \ldots W_{q-1}$ are pairwise disjoint and hence $w_{q-1}, w_1 \notin W_3$. So, there are no edges between W_0 and W_3, as required.

To establish the second bullet point, again use $k = 0$ and the fact that W_0, $\ldots, W_{q-1}$ are pairwise disjoint. We have $|(W_0 \cup W_2) \cap \Omega_1| = |W_0 \cap \Omega_1| + |W_2 \cap \Omega_1|$. By Claim B, each of these terms is odd, so their sum is even. The same argument applies to $(W_1 \cup W_3) \cap \Omega_2$.

(II) W is not entirely contained in one biconnected component. If this is true, then by Observation 9.20, there exists an obstruction B with cycle $C \in \mathrm{Cy}(B)$ such that, for some $k \in \{0, \ldots, q-1\}$, w_k and w_{k+1} are contained in C, w_{k+1} is an exit of B (in particular, an articulation point), and w_{k+2} and w_{k+3} are not contained in B.

Since $w_{k+2} \neq w_{k+4}$ by Lemma 9.21, it follows that no $v \in W_{k+3}$ is in B, which implies that there is no edge between W_k and W_{k+3}, as required.

For the second item, observe that W_k and W_{k+2} must be disjoint, as w_k and w_{k+1} are in the biconnected component B, but w_{k+2} is not. We further claim that W_{k+1} and W_{k+3} are disjoint. To see this, observe first that $W_{k+1} = \{w_{k+1}\}$ since w_{k+1} is the only common neighbour of w_k and w_{k+2} as otherwise w_{k+2} would be contained in B. Then we have already established that no $v \in W_{k+3}$ is in B, which implies $w_{k+1} \notin W_{k+3}$.

The fact that W_k and W_{k+2} are disjoint implies that $|(W_k \cup W_{k+2}) \cap \Omega_{k+1}| = (|W_k \cap \Omega_{k+1}| + |W_{k+2} \cap \Omega_{k+1}|)$. By Claim B, each of these terms is odd, so their sum is even. Using the fact that W_{k+1} and W_{k+3} are disjoint, the same is true for W_{k+1} and W_{k+3}.

■

We are finally able to invoke Lemma 5.6: Recall first, that $q \geq 3$ and $q \neq 4$ from the beginning of the proof. Recall that we identify the sets $\mathcal{C}_i$ of Lemma 5.6 with the sets W_i. Unless H has a hardness gadget (in which case we are finished) the following hold.

(L5.5.1) holds by Claim B.

(L5.5.2) and (L5.5.3) hold by Lemma 9.22.

(L5.5.4) is established by Lemma 9.23.

There is a k such that (L5.6.1) and (L5.6.2) hold by Claim C.

Consequently, all conditions are satisfied and we obtain a hardness gadget according to Lemma 5.6. □

9.5. Proof of the Main Theorem. We can now prove Theorem 1.4, which we restate for convenience.

THEOREM 1.4. *Let H be a simple graph whose involution-free reduction H^* is K_4-minor free. If H^* contains at most one vertex, then $\oplus\textsc{Hom}(H)$ can be solved in polynomial time. Otherwise, $\oplus\textsc{Hom}(H)$ is $\oplus$P-complete and, assuming the randomised Exponential Time Hypothesis, it cannot be solved in time $\exp(o(|G|))$.*

Proof. By Theorem 1.1, for every graph G, $|\hom(G \to H)| = |\hom(G \to H^*)| \bmod 2$. It is trivial to count homomorphisms to a graph with at most one vertex. Suppose that H^* has at least two vertices. Then it suffices to show that $\oplus\textsc{Hom}(H^*)$ is $\oplus$P-complete and that $\oplus\textsc{Hom}(H^*)$ cannot be solved in time $\exp(o(|G|))$, unless the rETH fails.

Since H^* is involution-free and contains at least 2 vertices, there is an involution-free connected component H' of H^* with at least 2 vertices as well: If H is disconnected, it has at least 2 connected components, and at least one of those two components cannot be a single vertex, as otherwise, we obtain a non-trivial involution by

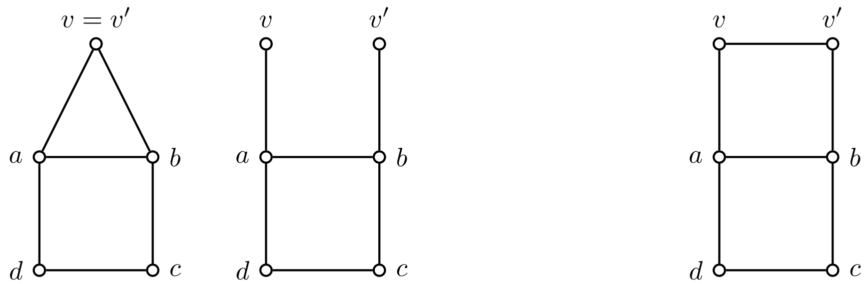

Fig. 12: Illustration of the two cases in the proof of Lemma 10.1.

switching those vertices. Furthermore, a connected component of an involution-free graph cannot have a non-trivial involution, as otherwise, the entire graph would have a non-trivial involution.

Next we claim that H' has a hardness gadget: Assume first that H' has a biconnected component that is not an edge, a diamond, an obstruction, or an impasse. By Lemma 8.10, H' has a hardness gadget. In the remaining case, every biconnected component of H' is an edge, a diamond, an obstruction, or an impasse. By Lemma 9.9, there is a closed suitable subtree T of the block-cut tree $\mathsf{BC}(H')$. If no block of T is an obstruction, then H' has a hardness gadget by Lemma 9.10. Otherwise, H' has a hardness gadget by Lemma 9.24.

This allows us to invoke Theorem 4.7 and we obtain that $\oplus\textsc{Ret}(H')$ is $\oplus$P-hard and cannot be solved in time $\exp(o|J|)$, unless the rETH fails.

Since H' is involution-free, we can reduce $\oplus\textsc{Ret}(H')$ to $\oplus\textsc{Hom}(H')$ by Theorem 4.4, and we can reduce $\oplus\textsc{Hom}(H')$ to $\oplus\textsc{Hom}(H^*)$ by Lemma 4.5. These reductions are tight in the sense that any subexponential-time algorithm for $\oplus\textsc{Hom}(H^*)$ would yield a subexponential-time algorithm for $\oplus\textsc{Ret}(H')$; this is due to the fact that the size of the oracle queries in each reduction is bounded linearly in the input size (see, e.g., the proof of Theorem 4.7 in which we made explicit that a linear bound on the size of the oracle queries is sufficient for the lower bound under rETH to transfer). We thus obtain $\oplus$P-hardness of $\oplus\textsc{Hom}(H^*)$, and that $\oplus\textsc{Hom}(H^*)$ cannot be solved in time $\exp(o(|G|))$, unless the rETH fails. $\square$

10. Counting Homomorphisms mod 2 to Graphs of Degree at most 3. We explored the possibilities for constructing hardness gadgets in graphs containing two squares that share one edge when we analysed K_4-minor-free and chordal bipartite graphs. It turns out that a similar strategy suffices to completely solve the case where H has degree at most 3. We start with the following lemma.

LEMMA 10.1. *Let H be an involution-free graph of degree at most 3 that contains a square. Then H has a hardness gadget.*

Proof. Let $C = (a, b, c, d, a)$ be a square in H. Assume first that at least one of the edges $\{a, c\}$ or $\{b, d\}$ are present. W.l.o.g. let $\{a, c\}$ be present. Then a and c have degree 3 and thus, by assumption, no further neighbours. Thus (ac) is a non-trivial involution of H.

Now assume that none of $\{a, c\}$ or $\{b, d\}$ are edges of H. If both, a and c have degree 2 then (ac) is an involution. Similarly, if b and d have both degree 2, we obtain the involution (bd). W.l.o.g. we can thus assume that a and b have degree 3. Let v

and v' be the neighbours of a and b, respectively, that are not contained in C. In what follows, we consider cases based on whether the edge $\{v, v'\}$ is present, and, if not, we differentiate between $v = v'$ and $v \neq v'$; an illustration is provided in Figure 12.

(I) $\{v, v'\} \notin E(H)$: This case corresponds to the two illustrations to the left of Figure 12. Note first that $\{v', d\}$ cannot be an edge of H, as otherwise, b and d both have neighbours $\{a, v', c\}$ (and no other neighbours, since they have degree 3), which means that (bd) is a non-trivial involution of H. Similarly, $\{v, c\}$ cannot be an edge of H, as otherwise (ac) is a non-trivial involution of H. Also, at least one of the edges $\{v, d\}$ and $\{v', c\}$ must not be present in H, as otherwise $(ad)(bc)$ is a non-trivial involution of H. W.l.o.g., assume that $\{v, d\}$ is not present. We construct a hardness gadget of H as follows:
   - $I = \{a\}$.
   - $S = \{b\}$.
   - J_1 is a path of 4 vertices: The first vertex is a b-pin, the second vertex is y, and the fourth vertex is an a-pin.
   - J_2 is a path of 3 vertices: The first vertex is an a-pin, the second vertex is z, and the third vertex is a c-pin.
   - J_3 is just the edge $\{y, z\}$.

   We first claim that $\Omega_y = \{v', a\}$. A vertex of H is in Ω_y if and only if it is adjacent to b and has an odd number of 2-paths to a. As H has degree at most three, the neighbours of b are precisely v', a and c. Note that a has degree precisely 3 and thus has an odd number of 2-paths to itself. Furthermore, there is only one 2-path from v' to a: This path contains b as internal vertex. There cannot be an additional 2-path from v' to a, since, in this case, the internal vertex must either be v, which is not possible as $\{v, v'\} \notin E(H)$, or d, which is not possible as $\{v, d\} \notin E(H)$. Finally, there are precisely two 2-paths from c to a: One has b as internal vertex, and the other has d as internal vertex. There cannot be a third one, as this 2-path would have v as internal vertex, but we ruled out the existence of the edge $\{v, c\}$. This shows that $\Omega_y = \{v', a\}$.
   Our next claim is that $\Omega_z = \{b, d\}$. Observe that Ω_z contains precisely the common neighbours of a and c. Thus b and d are included in Ω_z. The only candidate for a third common neighbour would be v, but we ruled out the existence of the edge $\{v, c\}$.
   Finally, we observe that $|\Sigma_{v',d}| = 0$ as $\{v', d\}$ is not an edge of H, and that $|\Sigma_{v',b}| = |\Sigma_{b,a}| = |\Sigma_{a,d}| = 1$.

(II) $\{v, v'\} \in E(H)$: This case corresponds to the illustration to the right of Figure 12. As in case (I), the edge $\{v, c\}$ is not present, as otherwise (ac) is a non-trivial involution, and that the edge $\{v', d\}$ is not present, as otherwise (bd) is a non-trivial involution. We construct a hardness gadget as follows:
   - $I = \{a\}$.
   - $S = \{b\}$.
   - J_1 is a path of 3 vertices: The first vertex is a v-pin, the second vertex is y, and the third vertex is a b-pin.
   - J_2 is a path of 3 vertices: The first vertex is an a-pin, the second vertex is z, and the third vertex is a c-pin.
   - J_3 is just the edge $\{y, z\}$.

   Note first that Ω_y contains precisely the common neighbours of v and b. Thus v' and a are contained in Ω_y. Recall further that c is not adjacent to v. As the degree of H is bounded by 3, we thus have $\Omega_y = \{v', a\}$. Similarly, we obtain that $\Omega_z = \{b, d\}$.

Finally, we observe that $|\Sigma_{v',d}| = 0$ as $\{v', d\}$ is not an edge of H, and that $|\Sigma_{v',b}| = |\Sigma_{b,a}| = |\Sigma_{a,d}| = 1$. ∎

Theorem 10.2. *Let H be a graph whose involution-free reduction H^* has maximum degree at most 3. If H^* contains at most one vertex, then $\oplus\textsc{Hom}(H)$ can be solved in polynomial time. Otherwise, $\oplus\textsc{Hom}(H)$ is $\oplus$P-complete and, assuming the randomised Exponential Time Hypothesis, it cannot be solved in time* $\exp(o(|G|))$.

Proof. By Theorem 1.1, for every graph G, $|\hom(G \to H)| = |\hom(G \to H^*)| \bmod 2$. It is trivial to count homomorphisms to a graph with at most one vertex. Suppose that H^* has at least two vertices. Then it suffices to show that $\oplus\textsc{Hom}(H^*)$ is $\oplus$P-complete and that $\oplus\textsc{Hom}(H^*)$ cannot be solved in time $\exp(o(|G|))$, unless the rETH fails.

If H^* does not contain a square but has at least 2 vertices, then it has a hardness gadget as shown in [19]. If H^* contains a square, then it has a hardness gadget by Lemma 10.1.

By Theorem 4.7, we obtain that $\oplus\textsc{Ret}(H^*)$ is $\oplus$P-hard and that it cannot be solved in time $\exp(o|J|)$, unless the rETH fails.

Finally, since H^* is involution-free, we can reduce $\oplus\textsc{Ret}(H^*)$ to $\oplus\textsc{Hom}(H^*)$ by Theorem 4.4. As we have already noted, the size of the oracle queries in this reduction are bounded linearly in the input size, so the reduction proves that any subexponential-time algorithm for $\oplus\textsc{Hom}(H^*)$ would yield a subexponential-time algorithm for $\oplus\textsc{Ret}(H^*)$, completing the proof. ∎

11. Counting List Homomorphisms modulo 2. Given graphs G and H together with a set of *lists* $\mathbf{S} = \{S_v \subseteq V(H) \mid v \in V(G)\}$, a *(list) homomorphism* from $(G, \mathbf{S})$ to H is a homomorphism h from G to H such that for each $v \in V(G)$ we have $h(v) \in S_v$. We use $\hom((G, \mathbf{S}) \to H)$ to denote the set of homomorphisms from $(G, \mathbf{S})$ to H. List homomorphisms are a natural generalisation of both homomorphisms and retractions.

In this section we provide a complete complexity classification for the problem of counting list homomorphisms modulo 2 to a given graph H. The classification determines for which graphs H the problem is feasible. We strengthen the result by considering a wider class of graphs H than in the rest of the paper (where we required H to be a simple graph, without self-loops or parallel edges). Let $\mathcal{H}$ be the set of all undirected graphs H which do not have parallel edges — self-loops are allowed.

Given a set S, let $\mathcal{P}(S)$ denote its power set. We consider the following problem, parameterised by a graph $H \in \mathcal{H}$ and by a set of lists $\mathcal{L} \subseteq \mathcal{P}(V(H))$.
Name: $\oplus\textsc{Hom}(H, \mathcal{L})$.
Input: A simple graph G and a collection of lists $\mathbf{S} = \{S_v \in \mathcal{L} \mid v \in V(G)\}$.
Output: $|\hom((G, \mathbf{S}) \to H)| \bmod 2$.

The input G to $\oplus\textsc{Hom}(H, \mathcal{L})$ is assumed to be simple because this is standard in the field, and because it makes results stronger. However, this restriction is not important for our result — see Remark 11.4. Taking $\mathcal{L} = P(V(H))$, the problem $\oplus\textsc{Hom}(H, \mathcal{P}(V(H)))$ is the problem of counting list homomorphisms to H modulo 2. To simplify the notation, we also write $\oplus\textsc{LHom}(H)$ for this problem. The following lemma is well-known.

Lemma 11.1. *Let H be a graph in $\mathcal{H}$ that contains a walk (a, b, c, d) such that $a \neq c$, $b \neq d$, and $\{a, d\} \notin E(H)$. Let $\mathcal{L} \subseteq \mathcal{P}(V(H))$ be a set of lists with $\{\{a, c\}, \{b, d\}\} \subseteq \mathcal{L}$. Then $\oplus\textsc{Hom}(H, \mathcal{L})$ is $\oplus$P-complete.*

Proof. The problem $\oplus$BIS, of counting the independent sets of a bipartite graph, modulo 2, is known to be $\oplus$P-complete [11, Theorem 4.2.1]. We will reduce $\oplus$BIS to $\oplus\textsc{Hom}(H,\mathcal{L})$.

Let G be a bipartite graph (an input to $\oplus$BIS) with vertex partition $V(G) = (L,R)$. For each $v \in L$, let $S_v = \{a,c\}$ and for each $v \in R$ let $S_v = \{b,d\}$. We set $\mathbf{S} = \{S_v \mid v \in V(G)\}$. Then every homomorphism h from $(G,\mathbf{S})$ to H corresponds to an independent set in G (and vice versa), where $h(v) \in \{a,d\}$ means that v is *in* the independent set and $h(v) \in \{b,c\}$ means that v is *out* of the independent set. (Since $a \neq c$ and $b \neq d$ it is well-defined whether v is in or out.) Hence a single $\oplus\textsc{LHom}(H,\mathcal{L})$ oracle call with input $(G,\mathbf{S})$ returns the number of independent sets of G, modulo 2. ∎

THEOREM 11.2. *Let H be a connected graph in $\mathcal{H}$ and let $\mathcal{L} \subseteq \mathcal{P}(V(H))$ be a set of lists with $\{S \subseteq V(H) \mid |S| = 2\} \subseteq \mathcal{L}$. If (i) H is a complete bipartite graph with no self-loops, or (ii) H is a complete graph in which every vertex has a self-loop, then $\oplus\textsc{Hom}(H,\mathcal{L})$ can be solved in polynomial time. Otherwise, $\oplus\textsc{Hom}(H,\mathcal{L})$ is $\oplus$P-complete.*

Proof. The easiness result comes from Dyer and Greenhill [9, Theorem 1.1]. (Dyer and Greenhill's result is stated for homomorphisms rather than for list homomorphisms, but it is easy to see, and well known, that it extends to list homomorphisms.) For the hardness part we consider four cases.

Case 1: H contains at least one looped and one unlooped vertex.

The problem $\oplus$IS, of counting the independent sets of a graph, modulo 2, is known to be $\oplus$P-complete [36]. In this case there is an easy reduction from $\oplus$IS to $\oplus\textsc{LHom}(H,\mathcal{L})$. To see this, note that, since H is connected, it contains a looped vertex a which is adjacent to an unlooped vertex b. Counting the homomorphisms from a graph G to $H[\{a,b\}]$ is well-known to be equivalent to counting the independent sets of G (see, e.g., [2]). Since $\{a,b\} \in \mathcal{L}$ we can use this list to restrict the image of homomorphisms to $\{a,b\}$, giving the desired reduction.

Case 2: H is a bipartite graph without self-loops but it is not a complete bipartite graph.

In this case, H contains a path (a,b,c,d) such that $\{a,d\} \notin E(H)$ so the problem $\oplus\textsc{Hom}(H,\mathcal{L})$ is $\oplus$P-complete by Lemma 11.1.

Case 3: H is a graph without self-loops that contains a cycle of odd length.

Consider a shortest odd-length cycle C in H. Due to minimality, C has to be an induced cycle of H (any additional edge between vertices of C would give a shorter even-length cycle and a shorter odd-length cycle). If C is not a triangle, then C contains a path (a,b,c,d) such that $\{a,d\} \notin E(H)$. If otherwise C is a triangle (a,b,c,a), then $\{a,a\} \notin E(H)$ since H does not have self-loops. In both cases $\oplus\textsc{Hom}(H,\mathcal{L})$ is $\oplus$P-complete by Lemma 11.1.

Case 4: H is a graph with all self-loops present but H is not a complete graph.

In this case, H contains a path (a,b,c) where $\{a,c\} \notin E(H)$. Since $\{b,b\} \in E(H)$ we can apply Lemma 11.1 to the walk (a,b,b,c) and obtain $\oplus$P-completeness of $\oplus\textsc{Hom}(H,\mathcal{L})$. ∎

The following complexity classification for the problem $\oplus\textsc{LHom}(H)$ follows easily from Theorem 11.2.

THEOREM 11.3. *Let H be graph in $\mathcal{H}$. If every connected component H' of H satisfies one of the following*

1. *H' is a complete bipartite graph with no self-loops, or*
2. *H' is a complete graph in which every vertex has a self-loop,*

then $\oplus$LHom(H) *can be solved in polynomial time. Otherwise,* $\oplus$LHom(H) *is* $\oplus$P-*complete.*

Proof. The easiness result comes from Dyer and Greenhill [9, Theorem 1.1]. For the hardness part, let H' be a connected component of H that is not a complete bipartite graph with no self-loops and is not a complete graph in which every vertex has a self-loop. Let $\mathcal{L}$ be the set of all size-2 subsets of $V(H')$. From Theorem 11.2, $\oplus$Hom$(H', \mathcal{L})$ is $\oplus$P-complete. However, $\oplus$Hom$(H', \mathcal{L})$ reduces trivially to $\oplus$LHom(H) — given an input $(G, \mathbf{S})$ to $\oplus$Hom$(H', \mathcal{L})$ simply return the number of (list) homomorphisms from $(G, \mathbf{S})$ to H, modulo 2.

Remark 11.4. Theorem 11.3 would be unchanged if we changed the definition of $\oplus$LHom(H) so that the input G can be any graph in $\mathcal{H}$ (so it need not be simple). A self-loop on a vertex v of G simply enforces the constraint that a homomorphism must map v to a vertex of H that has a self-loop. The same constraint can be enforced using a list.

Acknowledgements. We would like to thank Dave Richerby, our fellow member of the square-haters' club, for valuable discussions; this club will hopefully lose some members with the appearance of this work. Furthermore we thank Holger Dell for pointing out the tight conditional lower bound for counting independent sets modulo 2 in [5].

12. Index of Symbols and Definitions.

$H[S]$	subgraph of H induced by S		p. 8
homomorphism	edge-preserving mapping		p. 1
$\hom(G \to H)$	set of homomorphisms from G to H		p. 9
$\hom((J, x) \to (H, y))$	set of homomorphisms with pinned vertices		p. 9
impasses	distinguished family of chordal bipartite graphs	Def. 6.15	p. 23
involution	automorphism of order ≤ 2		p. 2
involution-free graph	graph without non-trivial involutions		p. 2
involution-free reduction			p. 3
K_4, K_4-minor-free	4-vertex complete graph		p. 3
list homomorphism			p. 62
$N_{W,H}(w_i)$	walk-neighbour-set	Def. 4.2	p. 9
obstruction	distinguished biconnected K_4-minor-free graph	Def. 8.6	p. 37
obstruction-free path	path in the block-cut tree excluding obstructions	Def. 9.11	p. 46
pair of connectors	distinguished pair of vertices of an impasse	Def. 6.15	p. 23
partially H-labelled graph	pair consisting of a graph and a pinning function		p. 9
$\mathrm{P}_H(a, b)$	a particular shortest path from a to b in H	Def. 9.12	p. 47
pinning function	partial function between vertices of two graphs		p. 9
pre-hardness gadget	$\oplus$P-hardness for all K_4-minor-free (1,2)-supergraphs	Def. 8.7	p. 37
R-closed/R-open		Def. 9.5	p. 44
rETH	randomised Exponential Time Hypothesis		p. 3
retraction	homomorphism from a partially labelled graph		p. 6
separation/separator		Def. 8.1	p. 35
$S_{k,\ell}$	distinguished V-typed supergraph of F	Def. 6.6	p. 16
strong hardness gadget	$\oplus$P-hardness for all K_4-minor-free supergraphs	Def. 6.8	p. 17
suitable connector		Def. 9.1	p. 42
suitable subtree		Def. 9.6	p. 44
type V	predicate for supergraphs of F	Def. 6.2	p. 15
walk-neighbour-set		Def. 4.2	p. 9
$\mathrm{W}_C(a, b)/\mathrm{W}_C(a)$	walk from a to b / from a to a along cycle C	Def. 9.16	p. 48
$\Gamma_H(v)$	neighbourhood of v in graph H		p. 9
$\Gamma_H(S)$	joint neighbourhood of S		p. 9
$\Gamma_{H \setminus F}(i, j)$	common neighbours of v_i and v_j in $H \setminus F$	Def. 6.1	p. 15
$\oplus\textsc{Hom}(H)$	counting homomorphisms to H mod 2		p. 2

$\oplus$IS	counting independent sets mod 2		p. 4
$\oplus$LHom(H)	counting list homomorphisms to H mod 2		p. 62
$\oplus$P	complexity class of parity problems		p. 2
$\oplus$Ret(H)	counting retractions to H mod 2		p. 6
(1,2)-supergraph	supergraph without new adjacencies and 2-paths	Def. 6.14	p. 23
$+$	concatenation of walks	Def. 9.15	p. 47

REFERENCES

[1] C. Borgs, J. Chayes, L. Lovász, V. T. Sós, and K. Vesztergombi, *Counting graph homomorphisms*, in Topics in Discrete Mathematics, M. Klazar, J. Kratochvíl, M. Loebl, J. Matoušek, P. Valtr, and R. Thomas, eds., Berlin, Heidelberg, 2006, Springer Berlin Heidelberg, pp. 315–371.

[2] G. R. Brightwell and P. Winkler, *Graph homomorphisms and phase transitions*, J. Comb. Theory, Ser. B, 77 (1999), pp. 221–262, https://doi.org/10.1006/jctb.1999.1899, https://doi.org/10.1006/jctb.1999.1899.

[3] H. Chen, R. Curticapean, and H. Dell, *The exponential-time complexity of counting (quantum) graph homomorphisms*, in Graph-Theoretic Concepts in Computer Science - 45th International Workshop, WG 2019, Vall de Núria, Spain, June 19-21, 2019, Revised Papers, I. Sau and D. M. Thilikos, eds., vol. 11789 of Lecture Notes in Computer Science, Springer, 2019, pp. 364–378, https://doi.org/10.1007/978-3-030-30786-8_28, https://doi.org/10.1007/978-3-030-30786-8_28.

[4] R. Curticapean, H. Dell, and D. Marx, *Homomorphisms are a good basis for counting small subgraphs*, in Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017, H. Hatami, P. McKenzie, and V. King, eds., ACM, 2017, pp. 210–223, https://doi.org/10.1145/3055399.3055502, https://doi.org/10.1145/3055399.3055502.

[5] H. Dell, T. Husfeldt, D. Marx, N. Taslaman, and M. Wahlen, *Exponential time complexity of the permanent and the Tutte polynomial*, ACM Trans. Algorithms, 10 (2014), pp. 21:1–21:32, https://doi.org/10.1145/2635812, https://doi.org/10.1145/2635812.

[6] J. Díaz, M. J. Serna, and D. M. Thilikos, *Counting H-colorings of partial k-trees*, Theor. Comput. Sci., 281 (2002), pp. 291–309, https://doi.org/10.1016/S0304-3975(02)00017-8, https://doi.org/10.1016/S0304-3975(02)00017-8.

[7] R. Diestel, *Graph Theory, 5th Edition*, vol. 173 of Graduate texts in mathematics, Springer, 2016.

[8] R. J. Duffin, *Topology of series-parallel networks*, Journal of Mathematical Analysis and Applications, 10 (1965), pp. 303–318.

[9] M. Dyer and C. Greenhill, *The complexity of counting graph homomorphisms*, Random Structures & Algorithms, 17 (2000), pp. 260–289.

[10] M. E. Dyer, L. A. Goldberg, C. S. Greenhill, and M. Jerrum, *The relative complexity of approximate counting problems*, Algorithmica, 38 (2004), pp. 471–500, https://doi.org/10.1007/s00453-003-1073-y, https://doi.org/10.1007/s00453-003-1073-y.

[11] J. Faben, *The complexity of counting solutions to generalised satisfiability problems modulo k*, arXiv preprint arXiv:0809.1836, (2008).

[12] J. Faben and M. Jerrum, *The Complexity of Parity Graph Homomorphism: An Initial Investigation*, Theory of Computing, 11 (2015), pp. 35–57, https://doi.org/10.4086/toc.2015.v011a002.

[13] T. Feder and P. Hell, *List homomorphisms to reflexive graphs*, J. Combin. Theory Ser. B, 72 (1998), pp. 236–250, https://doi.org/10.1006/jctb.1997.1812, https://ezproxy-prd.bodleian.ox.ac.uk:4563/10.1006/jctb.1997.1812.

[14] J. Flum and M. Grohe, *The parameterized complexity of counting problems*, SIAM J. Comput., 33 (2004), pp. 892–922, https://doi.org/10.1137/S0097539703427203, https://doi.org/10.1137/S0097539703427203.

[15] J. Focke, L. A. Goldberg, and S. Živný, *The complexity of approximately counting retractions to square-free graphs*, arXiv preprint arXiv:1907.02319, (2019).

[16] A. Galanis, L. A. Goldberg, and M. Jerrum, *Approximately counting H-colorings is #BIS-hard*, SIAM J. Comput., 45 (2016), pp. 680–711, https://doi.org/10.1137/15M1020551, https://doi.org/10.1137/15M1020551.

[17] A. Galanis, L. A. Goldberg, and M. Jerrum, *A complexity trichotomy for approximately counting list* H-*colorings*, ACM Trans. Comput. Theory, 9 (2017), pp. 9:1–9:22, https://doi.org/10.1145/3037381, https://doi.org/10.1145/3037381.

[18] A. Göbel, L. A. Goldberg, and D. Richerby, *The complexity of counting homomorphisms to cactus graphs modulo 2*, ACM Trans. Comput. Theory, 6 (2014), pp. 17:1–17:29, https://doi.org/10.1145/2635825, https://doi.org/10.1145/2635825.

[19] A. Göbel, L. A. Goldberg, and D. Richerby, *Counting homomorphisms to square-free graphs, modulo 2*, ACM Transactions on Computation Theory (TOCT), 8 (2016), p. 12.

[20] A. Göbel, J. A. G. Lagodzinski, and K. Seidel, *Counting homomorphisms to trees modulo a prime*, in 43rd International Symposium on Mathematical Foundations of Computer Science, MFCS 2018, August 27-31, 2018, Liverpool, UK, I. Potapov, P. G. Spirakis, and

J. Worrell, eds., vol. 117 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018, pp. 49:1–49:13, https://doi.org/10.4230/LIPIcs.MFCS.2018.49, https://doi.org/10.4230/LIPIcs.MFCS.2018.49.

[21] L. A. Goldberg and M. Jerrum, *The complexity of approximately counting tree homomorphisms*, ACM Trans. Comput. Theory, 6 (2014), pp. 8:1–8:31, https://doi.org/10.1145/2600917, https://doi.org/10.1145/2600917.

[22] L. M. Goldschlager and I. Parberry, *On the construction of parallel computers from various bases of Boolean functions*, Theor. Comput. Sci., 43 (1986), pp. 43–58.

[23] M. C. Golumbic and C. F. Goss, *Perfect Elimination and Chordal Bipartite Graphs*, Journal of Graph Theory, 2 (1978), pp. 155–163, https://doi.org/10.1002/jgt.3190020209, https://doi.org/10.1002/jgt.3190020209.

[24] M. Grohe, *The complexity of homomorphism and constraint satisfaction problems seen from the other side*, J. ACM, 54 (2007), pp. 1:1–1:24, https://doi.org/10.1145/1206035.1206036, https://doi.org/10.1145/1206035.1206036.

[25] M. Grohe, T. Schwentick, and L. Segoufin, *When is the evaluation of conjunctive queries tractable?*, in Proceedings on 33rd Annual ACM Symposium on Theory of Computing, July 6-8, 2001, Heraklion, Crete, Greece, J. S. Vitter, P. G. Spirakis, and M. Yannakakis, eds., ACM, 2001, pp. 657–666, https://doi.org/10.1145/380752.380867, https://doi.org/10.1145/380752.380867.

[26] P. Hell and J. Nesetril, *On the complexity of* H*-coloring*, J. Comb. Theory, Ser. B, 48 (1990), pp. 92–110, https://doi.org/10.1016/0095-8956(90)90132-J, https://doi.org/10.1016/0095-8956(90)90132-J.

[27] R. Impagliazzo and R. Paturi, *On the Complexity of k-SAT*, J. Comput. Syst. Sci., 62 (2001), pp. 367–375, https://doi.org/10.1006/jcss.2000.1727, https://doi.org/10.1006/jcss.2000.1727.

[28] 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, USA, R. E. Miller and J. W. Thatcher, eds., The IBM Research Symposia Series, Plenum Press, New York, 1972, pp. 85–103, https://doi.org/10.1007/978-1-4684-2001-2_9, https://doi.org/10.1007/978-1-4684-2001-2_9.

[29] A. Kazeminia and A. A. Bulatov, *Counting homomorphisms modulo a prime number*, in 44th International Symposium on Mathematical Foundations of Computer Science, MFCS 2019, August 26-30, 2019, Aachen, Germany, P. Rossmanith, P. Heggernes, and J. Katoen, eds., vol. 138 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019, pp. 59:1–59:13, https://doi.org/10.4230/LIPIcs.MFCS.2019.59, https://doi.org/10.4230/LIPIcs.MFCS.2019.59.

[30] S. Kelk, *On the relative complexity of approximately counting H-colourings*, PhD thesis, Warwick University, 2003.

[31] P. G. Kolaitis and M. Y. Vardi, *Conjunctive-query containment and constraint satisfaction*, J. Comput. Syst. Sci., 61 (2000), pp. 302–332, https://doi.org/10.1006/jcss.2000.1713, https://doi.org/10.1006/jcss.2000.1713.

[32] L. Lovász, *Graph minor theory*, Bull. Amer. Math. Soc. (N.S.), 43 (2006), pp. 75–86, https://doi.org/10.1090/S0273-0979-05-01088-8, https://doi.org/10.1090/S0273-0979-05-01088-8.

[33] C. H. Papadimitriou and S. Zachos, *Two remarks on the power of counting*, in Proc. 6th GI-Conference on Theoretical Computer Science, Springer-Verlag, 1982, pp. 269–275.

[34] M. Roth and P. Wellnitz, *Counting and finding homomorphisms is universal for parameterized complexity theory*, in Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020, S. Chawla, ed., SIAM, 2020, pp. 2161–2180, https://doi.org/10.1137/1.9781611975994.133, https://doi.org/10.1137/1.9781611975994.133.

[35] S. Toda, *PP is as hard as the polynomial-time hierarchy*, SIAM J. Comput., 20 (1991), pp. 865–877.

[36] L. G. Valiant, *Accidental algorthims*, in 2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06), IEEE, 2006, pp. 509–517.