

Counting Homomorphisms to Square-Free Graphs, Modulo 2

ANDREAS GÖBEL, LESLIE ANN GOLDBERG and DAVID RICHERBY, University of Oxford

We study the problem $\oplus\text{HOMSTO}H$ of counting, modulo 2, the homomorphisms from an input graph to a fixed undirected graph H . A characteristic feature of modular counting is that cancellations make wider classes of instances tractable than is the case for exact (non-modular) counting, so subtle dichotomy theorems can arise. We show the following dichotomy: for any H that contains no 4-cycles, $\oplus\text{HOMSTO}H$ is either in polynomial time or is $\oplus\text{P}$ -complete. This partially confirms a conjecture of Faben and Jerrum that was previously only known to hold for trees and for a restricted class of tree-width-2 graphs called cactus graphs. We confirm the conjecture for a rich class of graphs including graphs of unbounded tree-width. In particular, we focus on square-free graphs, which are graphs without 4-cycles. These graphs arise frequently in combinatorics, for example in connection with the strong perfect graph theorem and in certain graph algorithms. Previous dichotomy theorems required the graph to be tree-like so that tree-like decompositions could be exploited in the proof. We prove the conjecture for a much richer class of graphs by adopting a much more general approach.

Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—Computations on discrete structures; G.2.1 [Discrete Mathematics]: Combinatorics—Counting problems; G.2.2 [Discrete Mathematics]: Graph Theory—Graph algorithms

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Graph homomorphism, counting modulo 2, parity complexity dichotomy.

ACM Reference Format:

Andreas Göbel, Leslie Ann Goldberg, and David Richerby. 2015. Counting homomorphisms to square-free graphs modulo 2. *ACM Trans. Comput. Theory* V, N, Article A (January YYYY), 29 pages.
DOI: <http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

A homomorphism from a graph G to a graph H is a function from $V(G)$ to $V(H)$ that preserves edges, in the sense of mapping every edge of G to an edge of H ; non-edges of G may be mapped to edges or non-edges of H . Many structures arising in graph theory can be represented naturally as homomorphisms. For example, the proper q -colourings of a graph G correspond to the homomorphisms from G to a q -clique. For this reason, homomorphisms from G to a graph H are often called “ H -colourings” of G . Independent sets of G correspond to the homomorphisms from G to the connected graph with two vertices and one self-loop (vertices of G which are mapped to the self-loop are out of the corresponding independent set; vertices which are mapped to the other

Authors’ address: Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, Oxford OX1 3QD, UK.

The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007–2013) ERC grant agreement no. 334828. 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.

An extended abstract of this paper appeared in the proceedings of ICALP 2015.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 1942-3454/YYYY/01-ARTA \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>




Fig. 1. Theorem 1.2 shows that $\oplus\text{HOMSTO}H_1$ is $\oplus\text{P}$ -complete, whereas $\oplus\text{HOMSTO}H_2$ is in P. This, and the role of the starred vertex are explained later in the introduction.

vertex are in it). Homomorphism problems can also be seen as constraint satisfaction problems (CSPs) in which the constraint language consists of a single symmetric binary relation. Partition functions in statistical physics such as the Ising model, the Potts model, and the hard-core model arise naturally as weighted sums of homomorphisms [Bulatov and Grohe 2005; Goldberg et al. 2010].

In this paper, we study the complexity of counting homomorphisms modulo 2. For graphs G and H , $\text{Hom}(G \rightarrow H)$ denotes the set of homomorphisms from G to H . For each fixed H , we study the computational problem $\oplus\text{HOMSTO}H$, which is the problem of computing $|\text{Hom}(G \rightarrow H)| \bmod 2$, given an input graph G .

The structure of the graph H strongly influences the complexity of $\oplus\text{HOMSTO}H$. For example, consider the graphs H_1 and H_2 in Figure 1. Our result (Theorem 1.2) shows that $\oplus\text{HOMSTO}H_1$ is $\oplus\text{P}$ -complete, whereas $\oplus\text{HOMSTO}H_2$ is in P.

The aim of research in this area is to understand for which graphs H the problem $\oplus\text{HOMSTO}H$ is in P, for which graphs H the problem is $\oplus\text{P}$ -complete, and to prove that, for all graphs H , one or the other is true. Note that it isn't obvious, a priori, that there are no graphs H for which $\oplus\text{HOMSTO}H$ has intermediate complexity – proving that there are no such graphs H is the main work of a so-called *dichotomy theorem*.

This line of work was introduced by Faben and Jerrum [2015]. They made the following important conjecture (which requires a few definitions to state). An *involution* of a graph is an automorphism of order 2, i.e., an automorphism ρ that is not the identity but for which ρ^2 is the identity. Given a graph H and an involution ρ , H^ρ denotes the subgraph of H induced by the fixed points of ρ . We write $H \Rightarrow H'$ if there is an involution ρ of H such that $H^\rho = H'$ and we write $H \Rightarrow^* H'$ if either H is isomorphic to H' (written $H \cong H'$) or, for some positive integer k , there are graphs H_1, \dots, H_k such that $H \cong H_1$, $H_1 \Rightarrow \dots \Rightarrow H_k$, and $H_k \cong H'$. Faben and Jerrum showed [2015, Theorem 3.7] that for every graph H there is (up to isomorphism) exactly one involution-free graph H^* such that $H \Rightarrow^* H^*$. This graph H^* is called the *involution-free reduction* of H . See [Faben and Jerrum 2015, Figure 1] for a diagram showing a graph being reduced to its involution-free reduction. Faben and Jerrum make the following conjecture.

CONJECTURE 1.1. (*[Faben and Jerrum 2015]*) *Let H be a graph. If its involution-free reduction H^* has at most one vertex, then $\oplus\text{HOMSTO}H$ is in P; otherwise, $\oplus\text{HOMSTO}H$ is $\oplus\text{P}$ -complete.*

Note that our claim in Figure 1 is consistent with Conjecture 1.1. H_1 is involution-free, so it is its own involution-free reduction, but the involution-free reduction of H_2 is the single vertex marked $*$ in the figure.

Faben and Jerrum [2015, Theorem 3.8] proved Conjecture 1.1 for the case in which H is a tree. Subsequently, the present authors [Göbel et al. 2014, Theorem 1.6] proved the conjecture for a well-studied class of tree-width-2 graphs, namely *cactus graphs*, which are graphs in which each edge belongs to at most one cycle.

The main result of this paper is to prove the conjecture for a much richer class of graphs. In particular, we prove the conjecture for every graph H whose involution-free reduction has no 4-cycle (whether induced or not).

Graphs without 4-cycles are called “square-free” graphs. These graphs arise frequently in combinatorics, for example in connection with the strong perfect graph theorem [Conforti et al. 2004] and certain graph algorithms [Arends et al. 2011]. Our main theorem is the following.

THEOREM 1.2. *Let H be a graph whose involution-free reduction H^* is square-free. If H^* has at most one vertex, then $\oplus\text{HOMSTO}H$ is in P ; otherwise, $\oplus\text{HOMSTO}H$ is $\oplus\text{P}$ -complete.*

If H is square-free, then so is every induced subgraph, including its involution-free reduction H^* . Thus, we have the following corollary.

COROLLARY 1.3. *Let H be a square-free graph. If its involution-free reduction H^* has at most one vertex, then $\oplus\text{HOMSTO}H$ is in P ; otherwise, $\oplus\text{HOMSTO}H$ is $\oplus\text{P}$ -complete.*

In Section 1.3 we will discuss the reasons that we require H^* to be square-free in the proof of Theorem 1.2. First, in Section 1.1, we will describe the background to counting modulo 2. In Section 1.2, we will explain why Conjecture 1.1 is so much more difficult to prove for graphs with unbounded tree-width. Very briefly, in order to prove that $\oplus\text{HOMSTO}H$ is $\oplus\text{P}$ -hard without having a bound on the tree-width of H , it is necessary to take a much more abstract approach. Since it is not possible to decompose H using a tree-like decomposition as we did in [Göbel et al. 2014, Theorem 1.6], we have instead come up with an abstract characterisation of graph-theoretic structures in H which lead to $\oplus\text{P}$ -hardness. As we shall see, the proof that such structures always exist in square-free graphs involves interesting non-constructive elements, leading to a more abstract, and less technical (graph-theoretic) proof than [Göbel et al. 2014], while applying to a substantially richer set of graphs H , including graphs with unbounded tree width.

1.1. Counting modulo 2

Although counting modulo 2 produces a one-bit answer, the complexity of such problems has a rather different flavour from the complexity of decision problems. The complexity class $\oplus\text{P}$ was first studied by Papadimitriou and Zachos [1982] and by Goldschlager and Parberry [1986]. $\oplus\text{P}$ consists of all problems of the form “compute $f(x) \bmod 2$ ” where computing $f(x)$ is a problem in $\#\text{P}$. Toda [1991] has shown that there is a randomised polynomial-time reduction from every problem in the polynomial hierarchy to some problem in $\oplus\text{P}$. As such, $\oplus\text{P}$ is a large complexity class and $\oplus\text{P}$ -completeness seems to represent a high degree of intractability.

The unique flavour of modular counting is exhibited by Valiant’s famous restricted version of 3-SAT [Valiant 2006] for which counting solutions is $\#\text{P}$ -complete [Xia et al. 2007], counting solutions modulo 7 is in polynomial-time but counting solutions modulo 2 is $\oplus\text{P}$ -complete [Valiant 2006]. The seemingly mysterious number 7 was subsequently explained by Cai and Lu [2011], who showed that the k -SAT version of Valiant’s problem is tractable modulo any prime factor of $2^k - 1$.

Counting modulo 2 closely resembles ordinary, non-modular counting, but is still very different. Clearly, if a counting problem can be solved in polynomial time, the corresponding decision and parity problems are also tractable, but the converse does not necessarily hold. A characteristic feature of modular counting is cancellations, which can make the modular versions of hard counting problems tractable. For example, consider not-all-equal SAT, the problem of assigning values to Boolean variables such that

each of a given set of clauses contains both true and false literals. The number of solutions is always even, since solutions can be paired up by negating every variable in one solution to obtain a second solution. This makes counting modulo 2 trivial, while determining the exact number of solutions is $\#P$ -complete [Goldberg et al. 2014] and even deciding whether a solution exists is NP-complete [Schaefer 1978].

We use cancellations extensively in this paper. For example, if we wish to compute the size of a set S modulo 2 then, for any even-cardinality subset $X \subseteq S$, we have $|S| \equiv |S \setminus X| \pmod{2}$. This means that we can ignore the elements of X . It is also helpful to partition the set S into disjoint subsets S_1, \dots, S_ℓ exploiting the fact that $|S|$ is congruent modulo 2 to the number of odd-cardinality S_i . We use this idea frequently.

1.2. Going beyond bounded tree-width

1.2.1. Trees. All known hardness results for counting homomorphisms modulo 2 start with the following basic “pinning” approach. Let p be a function from $V(G)$ to $2^{V(H)}$. A homomorphism $f \in \text{Hom}(G \rightarrow H)$ *respects* the pinning function p if, for every $v \in V(G)$, $f(v)$ is in the set $p(v)$. Let $\text{PinHom}(G, H, p)$ be the set of homomorphisms from G to H that respect the pinning function p and let $\oplus\text{PINNEDHOMSTO}H$ be the problem of counting, modulo 2, the number of homomorphisms in $\text{PinHom}(G, H, p)$, given an input graph G and a pinning function p .

Faben and Jerrum [2015, Corollary 4.18] give a polynomial-time Turing reduction from the problem $\oplus\text{PINNEDHOMSTO}H$ to the problem $\oplus\text{HOMSTO}H$ for the special case in which the pinning function pins only two vertices of G , and these are both pinned to entire orbits of the automorphism group of H . The reduction relies on a result of Lovász [1967].

In order to use the reduction, it is necessary to show that the special case of the problem $\oplus\text{PINNEDHOMSTO}H$ is itself $\oplus P$ -hard. Faben and Jerrum restrict their attention to the case in which H is a tree, and this is helpful. Every involution-free tree is asymmetric (so the orbit of every vertex is trivial), so the pinning function p is actually able to pin two vertices of G to any two *particular* vertices of H .

The reduction that they used to prove hardness of $\oplus\text{PINNEDHOMSTO}H$ is from $\oplus\text{IS}$, the problem of counting independent sets modulo 2, which was shown to be $\oplus P$ -complete by Valiant [2006].

We first give an informal description of a general reduction from $\oplus\text{IS}$ to the problem $\oplus\text{PINNEDHOMSTO}H$. (The general description is actually based on our current approach in this paper, but we can also present past approaches in this context.) The vertices and edges of an input G of $\oplus\text{IS}$ are replaced by gadgets to give a graph J . In J , the gadget corresponding to the vertex v of G has a vertex y^v . We also choose an appropriate vertex i in H . Any homomorphism σ from J to the target graph H defines a set $I(\sigma) = \{v \in V(G) \mid \sigma(y^v) = i\}$ (mnemonic: “ i ” means “in” because $\sigma(y^v)$ is i exactly when v is in $I(\sigma)$). The configuration of the gadgets ensures that a set $I \subseteq V(G)$ has an odd number of homomorphisms σ with $I(\sigma) = I$ if and only if I is an independent set of G . Next, the homomorphisms $\sigma \in \text{Hom}(J \rightarrow H)$ can be partitioned according to the value of $I(\sigma)$. By the partitioning argument mentioned at the end of Section 1.1, the number of independent sets in G is equivalent to $|\text{Hom}(J \rightarrow H)|$, modulo 2.

The gadgets are chosen according to the structure and properties of H . Since Faben and Jerrum were working with trees, they were able to use gadgets with very simple structure: their gadgets are essentially paths and they exploit the fact that any non-trivial involution-free tree has at least two even-degree vertices and, of course, these have a unique path between them (which turns out to be useful).

1.2.2. Cactus graphs. The situation for cactus graphs is much more complicated. Non-trivial involution-free cactus graphs still contain even-degree vertices but the presence

of cycles means that paths, even shortest paths, are no longer guaranteed to be unique. Our solution in [Göbel et al. 2014] was to use more complicated gadgets. They are still (loosely) based on paths, since they are defined in terms of numbers of walks between vertices of H . However, rather than requiring appropriate even-degree vertices (which might not exist), we used a second, and more complicated, gadget to “select” an even-cardinality subset of a vertex’s neighbours. To find such gadgets in H , we used tree-like decompositions. Given a decomposition that breaks H into independent fragments, we inductively found gadgets (or, sometimes, partial gadgets) in the fragments, carefully putting them together across the join of the decomposition. All of this led to a very technical, very graph-theoretic solution, and also to a solution that does not generalise to graphs without tree-like decompositions.

The proof is complicated by the fact that there are involution-free graphs (even involution-free cactus graphs!) that have non-trivial automorphisms, unlike the situation for trees. Thus, the fact that the pinning function pins vertices to entire orbits (rather than to particular vertices) causes complications. The solution in [Göbel et al. 2014, Section 8] relies on special properties of cactus graphs, and it is not clear how it could be generalised.

1.2.3. Unbounded tree-width. Since they are based around a tree-like decomposition, the techniques of [Göbel et al. 2014] are not suitable for graphs with unbounded tree-width. To prove Conjecture 1.1 for a richer class of graphs, we adopt a much more abstract approach. Since we do not have tree-like decompositions, we instead mostly use structural properties of the whole graph to find gadgets. The structural properties do not always require technical detail – as we will see below, re-examining a result of Lovász [1967] even allows us to demonstrate non-constructively the existence of some of the gadgets that we use.

In order to support our more general approach, we first have to modify the pinning problem $\oplus\text{PINNEDHOMSTO}H$. For any graph H , a *partially H -labelled graph* $J = (G, \tau)$ consists of an *underlying graph* G and a *pinning function* τ , which in this paper is a partial function from $V(G)$ to $V(H)$. Thus, every vertex v in the domain of τ is pinned to a *particular* vertex of H and *not* to a subset such as an orbit. A homomorphism from a partially labelled graph $J = (G, \tau)$ to H is a homomorphism $\sigma: G \rightarrow H$ such that, for all vertices $v \in \text{dom}(\tau)$, $\sigma(v) = \tau(v)$. The intermediate problem that we study then is $\oplus\text{PARTLABHOMSTO}H$, the problem of computing $|\text{Hom}(J \rightarrow H)| \bmod 2$, given a partially H -labelled graph J . In Section 3, we generalise the application of Lovász’s theorem to show (Theorem 3.1) that $\oplus\text{PARTLABHOMSTO}H \leq \oplus\text{HOMSTO}H$.

Armed with a stronger pinning technique, we then abstract away most of the complications that arose for graphs with small tree-width by instead using more general gadgets, defined in Section 4. Because they are not based on paths, they do not rely on uniqueness of any path in H . Instead, the gadgets have three main parts. Our new reduction from $\oplus\text{IS}$ to $\oplus\text{HOMSTO}H$ can be seen informally as assigning colours to both the vertices and the edges of G , where each “colour” is a vertex of H . One part of the gadget controls which colours can be assigned to each vertex, one controls which colours can be assigned to each edge and a third part determines how many homomorphisms there are from G to H , given the choice of colours for the vertices and edges. In addition to all of this, we identify two special vertices of H , one of which is the vertex i mentioned above.

The much more general nature of our gadgets compared to those used previously makes them much easier to find and, in some cases, allows us to prove the existence of parts of them non-constructively. (Recall that gadgets depend only on the fixed graph H and not on the input G so they can be hard-coded into the reduction — there is no need to find one constructively.) We no longer need to find unique shortest paths in H or,

indeed, any paths at all. In fact, all the gadgets that we construct in this paper use a “caterpillar gadget” (Definition 4.3) which allows us to use *any* specified path in the graph H instead of relying on a unique shortest path. Rather than finding hardness gadgets in components in some decomposition of H , we mostly find gadgets “in situ”.

When a graph has two even-degree vertices, we can directly use those vertices and a caterpillar gadget to produce a hardness gadget (see Lemma 5.3). This already provides a self-contained proof of Faben and Jerrum’s dichotomy for trees. Next, for graphs with only one even-degree vertex, we show (Corollary 5.5) that deleting an appropriate set of vertices leaves a component with two even-degree vertices and show (Lemma 5.7) how to simulate that vertex deletion with gadgets. This leaves only graphs in which every vertex has odd degree. In such a graph, we are able to use any shortest odd-length cycle to construct a gadget (Lemma 5.13). If there are no odd cycles, the graph is bipartite. In this interesting case (Lemma 5.15) we use our version of Lovász’s result to find a gadget non-constructively.

1.3. Squares

It is natural to ask why the involution-free reduction H^* in Theorem 1.2 is required to be square-free. We do not believe that the restriction to square-free graphs is fundamental, since our results on pinning apply to all involution-free graphs (Section 3) and neither our definition of hardness gadgets (Definition 4.1) nor our proof that the existence of a hardness gadget for H implies that $\oplus\text{HOMSTO}H$ is $\oplus\text{P}$ -complete (Theorem 4.2) requires H to be square-free. However, all the actual hardness gadgets that we find for graphs do rely on the absence of 4-cycles, as discussed in Section 4.3, and removing this restriction seems technically challenging. We note that dealing with 4-cycles also caused significant difficulties in cactus graphs [Göbel et al. 2014].

1.4. Related work

We have already mentioned earlier work on counting graph homomorphisms modulo 2. The problem of counting graph homomorphisms (exactly, rather than modulo a fixed constant) was previously studied by Dyer and Greenhill [2000]. They showed the problem of counting homomorphisms to a fixed graph H is solvable in polynomial time if every connected component of H is a complete graph with a self-loop on every vertex or a complete bipartite graph with no self-loops, and is $\#\text{P}$ -complete, otherwise. Their work builds on an earlier dichotomy by Hell and Nešetřil [1990] for the complexity of the graph homomorphism decision problem (the problem of distinguishing between the case where there are no homomorphisms and the case where there is at least one). For work on counting modulo k in the *constraint satisfaction* setting see [Guo et al. 2011].

1.5. Organisation

We introduce notation in Section 2. Section 3 deals with pinning and consists mostly of adapting existing work to the precise framework we require. It can be skipped by the reader who is comfortable with pinning and happy to believe it can be done in our more general setting.

The gadgets that we use are formally defined in Section 4, where we also show that $\oplus\text{HOMSTO}H$ is $\oplus\text{P}$ -complete if H is an involution-free graph that has one of these gadgets. Section 4.2 introduces a gadget that we use extensively, but which requires H to be square-free, as discussed in Section 4.3. In Section 5, we show how to find hardness gadgets for all square-free graphs and, in Section 6, we tie everything together to prove the dichotomy theorem.

2. NOTATION

We write $[n]$ for the set $\{1, \dots, n\}$. For a set S and an element x , we often write $S - x$ for $S \setminus \{x\}$.

Graphs. In this paper, graphs are undirected and have no parallel edges and no loops. The one exception to this is that we briefly allow loops in the proof of Lemma 3.6 (this is clearly stated in the proof). Paths and cycles do not repeat vertices; walks may repeat both vertices and edges. The length of a path or cycle is the number of edges that it contains. The *odd-girth* of a graph is the length of its shortest odd-length cycle. $\Gamma_G(v)$ is the set of neighbours of a vertex v in G .

We write $G \cong H$ to indicate that graphs G and H are isomorphic. $\text{Aut}(H)$ denotes the automorphism group of a graph H . An *involution* is an automorphism of order 2 (i.e., an automorphism ρ that is not the identity such that $\rho \circ \rho$ is the identity). $\text{Hom}(G \rightarrow H)$ denotes the set of homomorphisms from a graph G to a graph H .

Partially labelled graphs. For any graph H , a *partially H -labelled graph* $J = (G, \tau)$ consists of an *underlying graph* G and a *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* or *pinned to* $\tau(v)$. We will refer to these graphs as *partially labelled graphs* where the graph H is clear from the context. We sometimes write $G(J)$ and $\tau(J)$ for the underlying graph and pinning function of a partially labelled graph, respectively. We write partial functions as sets of pairs, for example, writing $\tau = \{a \mapsto s, b \mapsto t\}$ for the partial function τ with $\text{dom}(\tau) = \{a, b\}$ such that $\tau(a) = s$ and $\tau(b) = t$.

A homomorphism from a partially labelled graph $J = (G, \tau)$ to H is a homomorphism $\sigma: G \rightarrow H$ such that, for all vertices $v \in \text{dom}(\tau)$, $\sigma(v) = \tau(v)$. We say that such a homomorphism *respects* τ .

Distinguished vertices. It is often convenient to regard a graph as having some number of distinguished vertices x_1, \dots, x_r and we denote such a graph by (G, x_1, \dots, x_r) . Note that the distinguished vertices need not be distinct. We sometimes abbreviate the sequence x_1, \dots, x_r as \bar{x} and we use $G[\bar{x}]$ to denote the subgraph of G induced by the set of vertices $\{x_1, \dots, x_r\}$. A homomorphism from a graph (G, x_1, \dots, x_r) to (H, y_1, \dots, y_r) is a homomorphism σ from G to H with the property that $\sigma(x_i) = y_i$ for each $i \in [r]$. This is the same thing as a homomorphism from the partially H -labelled graph $(G, \{x_1 \mapsto y_1, \dots, x_r \mapsto y_r\})$ to H . Given a partially labelled graph $J = (G, \tau)$ and vertices $x_1, \dots, x_r \notin \text{dom}(\tau)$, a homomorphism from (J, x_1, \dots, x_r) to (H, y_1, \dots, y_r) is formally identical to a homomorphism from $J' = (G, \tau \cup \{x_1 \mapsto y_1, \dots, x_r \mapsto y_r\})$ to H .

Similarly, we say that two graphs (G, x_1, \dots, x_r) and (H, y_1, \dots, y_s) are isomorphic if $r = s$ and there is an isomorphism $\rho: V(G) \rightarrow V(H)$ such that $\rho(x_i) = y_i$ for each $i \in [r]$ (note that we may have $G = H$). An automorphism of (G, x_1, \dots, x_r) is just an automorphism ρ of G with the property that $\rho(x_i) = x_i$ for each $i \in [r]$.

Diagram conventions. In diagrams of partially labelled graphs, ordinary vertices are denoted by black dots, distinguished vertices by small white circles and pinned vertices (i.e., the vertices in $\text{dom}(\tau)$) by large white circles. A label next to a vertex of any kind indicates the identity of that vertex; a label inside a white circle indicates what that vertex is pinned to.

3. PARTIALLY LABELLED GRAPHS AND PINNING

The results in this section do not require H to be square-free.

We use pinning in our gadgets, so we mostly work with the problem of determining the number of homomorphisms from a partially H -labelled graph to H , modulo 2:

Name. $\oplus\text{PARTLABHOMSTO}H$.

Parameter. A graph H .

Input. A partially H -labelled graph J .

Output. $|\text{Hom}(J \rightarrow H)| \bmod 2$.

Our goal in the remainder of this section is to prove the following theorem.

THEOREM 3.1. $\oplus\text{PARTLABHOMSTOH} \leq \oplus\text{HOMSTOH}$ for any involution-free graph H .

The reader who is prepared to take Theorem 3.1 on trust may safely skip the rest of this section. The theorem is used in later sections but the details of its proof are not.

To prove the theorem, we need to develop some machinery. This closely follows the presentation of similar material by Faben and Jerrum [2015] and our earlier paper [Göbel et al. 2014] which, in turn, draw on the work of Lovász [1967] and Hell and Nešetřil [2004]. This duplication is unfortunate but, at the end of the section, we explain how the results we have presented are subtly different from those in the literature so existing results could not be reused directly.

After stating some elementary group theory results that we need, we prove in Section 3.2 a version of a result originally due to Lovász. This (Lemma 3.6) states that, if graphs with distinguished vertices (H, \bar{y}) and (H', \bar{y}') are non-isomorphic, there is a graph (G, \bar{x}) that has an odd number of homomorphisms to one of (H, \bar{y}) and (H', \bar{y}') and an even number of homomorphisms to the other. Taking $H' = H$, this allows us to distinguish two tuples of vertices in H from one another, as long as they are not in the same orbit of $\text{Aut}(H)$.

This is not quite enough for pinning, as it doesn't give us control over which of the two graphs receives an odd number of homomorphisms from (G, \bar{x}) . In Section 3.3, we solve this problem algebraically, adapting a technique of Faben and Jerrum [2015]. This allows us to prove Theorem 3.1 in Section 3.4 and thereby implement the pinning we need for our reductions.

3.1. Group-theoretic background

We will require two results from group theory. For the first, see, e.g., [Armstrong 1988, Theorem 13.1].

THEOREM 3.2 (CAUCHY'S GROUP THEOREM). *If \mathcal{G} is a finite group and a prime p divides $|\mathcal{G}|$, then \mathcal{G} contains an element of order p .*

For a permutation group \mathcal{G} acting on a set X , the *orbit* of an element $x \in X$ is the set $\text{Orb}_{\mathcal{G}}(x) = \{\pi(x) \mid \pi \in \mathcal{G}\}$. For a graph H , we will abuse notation mildly by writing $\text{Orb}_H(\cdot)$ instead of $\text{Orb}_{\text{Aut}H}(\cdot)$.

The following is a corollary of the orbit–stabiliser theorem [Armstrong 1988, Corollary 17.3].

THEOREM 3.3. *Let \mathcal{G} be a finite permutation group acting on a set X . For every $x \in X$, $|\text{Orb}_{\mathcal{G}}(x)|$ divides $|\mathcal{G}|$.*

These two theorems have the following corollary about the size of orbits under the automorphism group of involution-free graphs.

COROLLARY 3.4. *Let H be an involution-free graph. Every orbit of a tuple $\bar{y} \in V(H)^r$ under the action of $\text{Aut}(H)$ has odd cardinality.*

PROOF. By Theorem 3.2, $|\text{Aut}(H)|$ is odd, since the group contains no element of order 2. Consider the natural action of $\text{Aut}(H)$ on $V(H)^r$. By Theorem 3.3, the size of the orbit of \bar{y} in H divides $|\text{Aut}(H)|$ so is also odd. \square

3.2. A Lovász-style lemma

Lovász proved that two graphs H and H' are isomorphic if and only if $|\text{Hom}(G \rightarrow H)| = |\text{Hom}(G \rightarrow H')|$ for every graph G (in fact, he proved the analogous result for general relational structures but we do not need this here). We show that this result remains true even if we replace equality of the number of homomorphisms with equivalence modulo 2. Faben and Jerrum also showed this [2015, Lemma 3.13] though in a less general setting than the one that we need. Our proof is based on the presentation of Hell and Nešetřil [2004, Section 2.3].

For the proof we need some definitions, which are used only in this section. We say that two r -tuples \bar{x} and \bar{y} *have the same equality type* if, for all $i, j \in [r]$, $x_i = x_j$ if and only if $y_i = y_j$. Let $\text{InjHom}((G, \bar{x}) \rightarrow (H, \bar{y}))$ be the set of injective homomorphisms from (G, \bar{x}) to (H, \bar{y}) .

Before proving the main lemma, we prove a simple fact about injective homomorphisms and equality types of distinguished variables.

LEMMA 3.5. *Let (G, \bar{x}) and (H, \bar{y}) be graphs, each with r distinguished vertices. If \bar{x} and \bar{y} do not have the same equality type, then $|\text{InjHom}((G, \bar{x}) \rightarrow (H, \bar{y}))| = 0$.*

PROOF. If there are $i, j \in [r]$ such that $x_i = x_j$ but $y_i \neq y_j$, then there are no homomorphisms (injective or otherwise) from (G, \bar{x}) to (H, \bar{y}) , since x_i cannot be mapped simultaneously to both y_i and y_j . Otherwise, there must be $i, j \in [r]$ such that $x_i \neq x_j$ but $y_i = y_j$. Then no homomorphism η can be injective because we must have $\eta(x_i) = \eta(x_j) = y_i$. \square

LEMMA 3.6. *Let (H, \bar{y}) and (H', \bar{y}') be involution-free graphs, each with r distinguished vertices. Then $(H, \bar{y}) \cong (H', \bar{y}')$ if and only if, for all (not necessarily connected) graphs (G, \bar{x}) with r distinguished vertices,*

$$|\text{Hom}((G, \bar{x}) \rightarrow (H, \bar{y}))| \equiv |\text{Hom}((G, \bar{x}) \rightarrow (H', \bar{y}'))| \pmod{2}. \quad (1)$$

PROOF. If (H, \bar{y}) and (H', \bar{y}') are isomorphic, it follows trivially that (1) holds for all graphs (G, \bar{x}) . For the other direction, suppose that (1) holds for all (G, \bar{x}) .

First, we claim that this implies that \bar{y} and \bar{y}' have the same equality type. If they have different equality types then, without loss of generality, we may assume that there are distinct indices i and j such that $y_i = y_j$ but $y'_i \neq y'_j$. Let G be the graph on vertices $\{y_1, \dots, y_r\}$ with no edges: we see that $|\text{Hom}((G, \bar{y}) \rightarrow (H, \bar{y}))| = 1 \neq |\text{Hom}((G, \bar{y}) \rightarrow (H', \bar{y}'))| = 0$, contradicting the assumption that (1) holds for all G .

Second, we show by induction on the number of vertices in G that, if (1) holds for all (G, \bar{x}) then, for all (G, \bar{x}) ,

$$|\text{InjHom}((G, \bar{x}) \rightarrow (H, \bar{y}))| \equiv |\text{InjHom}((G, \bar{x}) \rightarrow (H', \bar{y}'))| \pmod{2}, \quad (2)$$

Specifically, under the assumption that (1) holds for all (G, \bar{x}) , we show that (2) holds for all (G, \bar{x}) with $|V(G)| \leq n_0$ for a suitable value n_0 and that, if (2) holds for all (G, \bar{x}) with $|V(G)| < n$, it also holds for any (G, \bar{x}) with $|V(G)| = n$.

Let $n_0 = |\{y_1, \dots, y_r\}| = |\{y'_1, \dots, y'_r\}|$ be the number of distinct elements in \bar{y} . For the base case of the induction, consider any graph (G, \bar{x}) with $|V(G)| \leq n_0$. If \bar{x} does not have the same equality type as \bar{y} and \bar{y}' (which is guaranteed if $|V(G)| < n_0$) then, by Lemma 3.5,

$$|\text{InjHom}((G, \bar{x}) \rightarrow (H, \bar{y}))| = |\text{InjHom}((G, \bar{x}) \rightarrow (H', \bar{y}'))| = 0.$$

If \bar{x} has the same equality type as \bar{y} and \bar{y}' then, in particular, every vertex of G is distinguished. Any homomorphism from (G, \bar{x}) to (H, \bar{y}) or (H', \bar{y}') is injective, so

$$\begin{aligned} |\text{InjHom}((G, \bar{x}) \rightarrow (H, \bar{y}))| &= |\text{Hom}((G, \bar{x}) \rightarrow (H, \bar{y}))| \\ &= |\text{Hom}((G, \bar{x}) \rightarrow (H', \bar{y}'))| \\ &= |\text{InjHom}((G, \bar{x}) \rightarrow (H', \bar{y}'))|, \end{aligned}$$

where the second equality is by the assumption that (1) holds for (G, \bar{x}) .

For the inductive step, let $n > n_0$ and assume that (2) holds for all (G, \bar{x}) with $|V(G)| < n$. Now, consider some (G, \bar{x}) with $|V(G)| = n$.

Given any homomorphism σ from (G, \bar{x}) to (H, \bar{y}) , we can define an equivalence relation θ on $V(G)$ by $(u, v) \in \theta$ if and only if $\sigma(u) = \sigma(v)$. (Note that, if σ is injective, then θ is just the equality relation on $V(G)$.) Write $[u]$ for the θ -equivalence class of a vertex $u \in V(G)$. Let G/θ be the graph whose vertex set is $\{[u] \mid u \in V(G)\}$ and whose edge set is $\{([u], [v]) \mid (u, v) \in E(G)\}$. For graphs with distinguished vertices, we write $(G, x_1, \dots, x_r)/\theta = (G/\theta, [x_1], \dots, [x_r])$. The homomorphism σ from (G, \bar{x}) to (H, \bar{y}) corresponds to an injective homomorphism from $(G, \bar{x})/\theta$ to (H, \bar{y}) .

Note that, if there are adjacent vertices u and v in G such that $(u, v) \in \theta$ for some equivalence relation θ , the graph G/θ has a self-loop on the vertex $[u]$. This is not a problem. Because H is loop-free, there are no homomorphisms (injective or otherwise) from such a graph G/θ to H . For the same reason, there are no homomorphisms from G to H that map adjacent vertices u and v to the same place. Therefore, this particular θ does not correspond to any homomorphism from G to H and contributes zero to the sums below, as required.

We have

$$\begin{aligned} |\text{Hom}((G, \bar{x}) \rightarrow (H, \bar{y}))| &= |\text{InjHom}((G, \bar{x}) \rightarrow (H, \bar{y}))| + \sum_{\theta} |\text{InjHom}((G, \bar{x})/\theta \rightarrow (H, \bar{y}))| \\ |\text{Hom}((G, \bar{x}) \rightarrow (H', \bar{y}'))| &= |\text{InjHom}((G, \bar{x}) \rightarrow (H', \bar{y}'))| + \sum_{\theta} |\text{InjHom}((G, \bar{x})/\theta \rightarrow (H', \bar{y}'))|, \end{aligned}$$

where the sums are over all equivalence relations θ , except for the equality relation.

The left-hand sides of these equations are equivalent modulo 2 by assumption. The sums over θ on the right are equivalent modulo 2 by the inductive hypothesis since θ is not the equality relation, so G/θ has fewer vertices than G . Therefore, (2) holds for the graph under consideration.

Finally, it remains to prove that (2) holding for all (G, \bar{x}) implies that $(H, \bar{y}) \cong (H', \bar{y}')$. To see this, take $(G, \bar{x}) = (H, \bar{y})$. An injective homomorphism from a graph to itself is an automorphism and, since (H, \bar{y}) is involution-free, $\text{Aut}(H, \bar{y})$ has no element of order 2, so $|\text{Aut}(H, \bar{y})|$ is odd by Cauchy's group theorem (Theorem 3.2). By (2), there are an odd number of injective homomorphisms from (H, \bar{y}) to (H', \bar{y}') , which means that there is at least one such homomorphism. Similarly, taking $(G, \bar{x}) = (H', \bar{y}')$ shows that there is an injective homomorphism from (H', \bar{y}') to (H, \bar{y}) and, therefore, the two graphs are isomorphic. \square

For our nonconstructive proof that some gadgets exist, we use the following corollary of the proof of Lemma 3.6, which restricts to a certain class of connected graphs.

COROLLARY 3.7. *Let (H, \bar{y}) and (H', \bar{y}') be connected, involution-free graphs, each with r distinguished vertices, such that $H[\bar{y}]$ and $H'[\bar{y}']$ are also connected. Then $(H, \bar{y}) \cong (H', \bar{y}')$ if and only if (1) holds for all connected graphs (G, \bar{x}) with r distinguished vertices such that $G[\bar{x}]$ is connected.*

PROOF. For brevity, we refer to (G, \bar{x}) as *appropriate* if it is connected, it has r distinguished vertices and $G[\bar{x}]$ is connected.

As in the proof of Lemma 3.6, the “only if” direction is trivial, so we suppose that (1) holds for all appropriate (G, \bar{x}) . Also, \bar{y} and \bar{y}' must have the same equality type. If they do not, we may assume there are distinct i and j with $y_i = y_j$ but $y'_i \neq y'_j$, and take $G = H[\bar{y}]$. (G, \bar{y}) is appropriate but we have $|\text{Hom}((G, \bar{y}) \rightarrow (H, \bar{y}))| = 1 \neq |\text{Hom}((G, \bar{y}) \rightarrow (H, \bar{y}'))| = 0$, which contradicts the assumption that (1) holds for all appropriate (G, \bar{x}) .

The proof that (1) holding for every appropriate G implies that (2) holds for every appropriate G proceeds by induction on $|V(G)|$, as in the proof of the lemma. The base cases are unchanged. To see that the inductive step remains valid, let (G, \bar{x}) be appropriate and let θ be any equivalence relation on $V(G)$. We claim that $(G, \bar{x})/\theta$ is also appropriate. By construction, $(G, \bar{x})/\theta$ has r distinguished vertices. It is connected because it is the result of identifying vertices in a connected graph; $(G/\theta)[[x_1], \dots, [x_r]]$ is connected for the same reason.

This establishes that (2) holds for all appropriate (G, \bar{x}) . Since (H, \bar{y}) and (H', \bar{y}') are both appropriate, we can complete the proof in the same way as in the proof of Lemma 3.6, substituting each of these graphs in turn for (G, \bar{x}) in (2). \square

3.3. Implementing vectors

The presentation in this section follows very closely that of Faben and Jerrum [2015], extended to r -tuples of distinguished vertices.

Definition 3.8. Let H be an involution-free graph. We refer to a list $\bar{y}_1, \dots, \bar{y}_\lambda$ of elements of $V(H)^r$ as an *enumeration of $V(H)^r$ up to isomorphism* if, for every $\bar{y} \in V(H)^r$, there is exactly one $i \in [\lambda]$ such that $(H, \bar{y}) \cong (H, \bar{y}_i)$.

Note that the number λ of tuples in the enumeration depends on H .

Definition 3.9. Let (G, \bar{x}) be a graph with r distinguished vertices. We define the vector $\mathbf{v}_H(G, \bar{x}) \in \{0, 1\}^\lambda$ where, for each $i \in [\lambda]$, the i th component of $\mathbf{v}_H(G, \bar{x})$ is given by

$$(\mathbf{v}_H(G, \bar{x}))_i \equiv |\text{Hom}((G, \bar{x}) \rightarrow (H, \bar{y}_i))| \pmod{2}.$$

We say that (G, \bar{x}) *implements* this vector.

Define \oplus and \otimes to be, respectively, component-wise addition and multiplication, modulo 2, of vectors in $\{0, 1\}^\lambda$.

LEMMA 3.10. *Let $\bar{x} = x_1 \dots x_r$ and let (G_1, \bar{x}) and (G_2, \bar{x}) be graphs such that $V(G_1) \cap V(G_2) = \{x_1, \dots, x_r\}$. Then,*

$$\mathbf{v}_H(G_1 \cup G_2, \bar{x}) = \mathbf{v}_H(G_1, \bar{x}) \otimes \mathbf{v}_H(G_2, \bar{x}).$$

PROOF. A function $\sigma: V(G_1) \cup V(G_2) \rightarrow V(H)$ is a homomorphism from $(G_1 \cup G_2, \bar{x})$ to (H, \bar{y}) if and only if, for each $i \in \{1, 2\}$, the restriction of σ to $V(G_i)$ is a homomorphism from (G_i, \bar{x}) to (H, \bar{y}) . \square

In contrast, given (G_1, \bar{x}_1) and (G_2, \bar{x}_2) , it is not obvious that there is a graph (G, \bar{x}) such that $\mathbf{v}_H(G, \bar{x}) = \mathbf{v}_H(G_1, \bar{x}_1) \oplus \mathbf{v}_H(G_2, \bar{x}_2)$. Following Faben and Jerrum [2015], we side-step this issue by introducing a formal sum of graphs. Given graphs with distinguished vertices $(G_1, \bar{x}_1), \dots, (G_t, \bar{x}_t)$, we define

$$\mathbf{v}_H((G_1, \bar{x}_1) + \dots + (G_t, \bar{x}_t)) = \mathbf{v}_H(G_1, \bar{x}_1) \oplus \dots \oplus \mathbf{v}_H(G_t, \bar{x}_t)$$

and we say that a vector $\mathbf{v} \in \{0, 1\}^\lambda$ is *H-implementable* if it can be expressed as such a sum.

We require the following, which is essentially Lemma 4.16 of [Faben and Jerrum 2015].

LEMMA 3.11. *Let $S \subseteq \{0, 1\}^\lambda$ be closed under \oplus and \otimes . If $1^\lambda \in S$ and, for every distinct $i, j \in [\lambda]$, there is a tuple $s = s_1 \dots s_\lambda \in S$ with $s_i \neq s_j$, then $S = \{0, 1\}^\lambda$.*

COROLLARY 3.12. *Let H be an involution-free graph. Every $\mathbf{v} \in \{0, 1\}^\lambda$ is H -implementable.*

PROOF. Let S be the set of H -implementable vectors. S is clearly closed under \oplus , and is closed under \otimes by Lemma 3.10. Let G be the graph on vertices $\{x_1, \dots, x_r\}$, with no edges. 1^λ is implemented by (G, x_1, \dots, x_r) , which has exactly one homomorphism to every (H, \bar{y}_i) . Finally, for every distinct pair $i, j \in [\lambda]$, (H, \bar{y}_i) and (H, \bar{y}_j) are not isomorphic, by definition of the enumeration of r -tuples (up to isomorphism). Therefore, by Lemma 3.6, there is a graph (G, \bar{x}) such that

$$|\text{Hom}((G, \bar{x}) \rightarrow (H, \bar{y}_i))| \not\equiv |\text{Hom}((G, \bar{x}) \rightarrow (H, \bar{y}_j))| \pmod{2}.$$

(G, \bar{x}) implements a vector \mathbf{v} whose i th and j th components are different. \square

3.4. Pinning

We now have almost everything we need to prove Theorem 3.1. Recall the definition of an enumeration $\bar{y}_1, \dots, \bar{y}_\lambda$ of $V(H)^r$ up to isomorphism (Definition 3.8).

LEMMA 3.13. *Let H be an involution-free graph and let $\bar{y}_1, \dots, \bar{y}_\lambda$ be an enumeration of $V(H)^r$ up to isomorphism. For any graph (G, \bar{x}) with r distinguished vertices,*

$$|\text{Hom}(G \rightarrow H)| \equiv \sum_{i \in [\lambda]} (\mathbf{v}_H(G, \bar{x}))_i \pmod{2}.$$

PROOF. We have (for details see below),

$$\begin{aligned} \sum_{i \in [\lambda]} (\mathbf{v}_H(G, \bar{x}))_i &\equiv \sum_{i \in [\lambda]} |\text{Hom}((G, \bar{x}) \rightarrow (H, \bar{y}_i))| \pmod{2} \\ &\equiv \sum_{i \in [\lambda]} |\text{Orb}_H(\bar{y}_i)| |\text{Hom}((G, \bar{x}) \rightarrow (H, \bar{y}_i))| \pmod{2} \\ &= \sum_{i \in [\lambda]} \sum_{\bar{y} \in \text{Orb}_H(\bar{y}_i)} |\text{Hom}((G, \bar{x}) \rightarrow (H, \bar{y}))| \\ &= |\text{Hom}(G \rightarrow H)|. \end{aligned}$$

The second equivalence modulo 2 is because all orbits have odd cardinality by Corollary 3.4 and multiplying the terms of the sum by odd numbers doesn't change the total, modulo 2. The first equality is because, for any $\bar{y} \in \text{Orb}_H(\bar{y}_i)$, $|\text{Hom}((G, \bar{x}) \rightarrow (H, \bar{y}))| = |\text{Hom}((G, \bar{x}) \rightarrow (H, \bar{y}_i))|$. This is because composing a homomorphism from (G, \bar{x}) to (H, \bar{y}) with an isomorphism from (H, \bar{y}) to (H, \bar{y}_i) gives a homomorphism from (G, \bar{x}) to (H, \bar{y}_i) . The final equality is because every homomorphism from G to H must map \bar{x} to some tuple \bar{y} and (exactly) all such tuples are included exactly once in the double sum. \square

We can now prove Theorem 3.1: for any involution-free graph H , there is a polynomial-time Turing reduction from $\oplus\text{PARTLABHOMSTO}H$ to $\oplus\text{HOMSTO}H$.

PROOF OF THEOREM 3.1. Let $J = (G, \tau)$ be an instance of $\oplus\text{PARTLABHOMSTO}H$. Let $\bar{x} = x_1, \dots, x_r$ be an enumeration of $\text{dom}(\tau)$ and let $\bar{y} = y_1, \dots, y_r = \tau(x_1), \dots, \tau(x_r)$.




Fig. 2. An involution-free graph H illustrating the difference between pinning vertices to orbits of vertices and pinning a tuple of vertices to an orbit of a tuple.

Moving from the world of partially H -labelled graphs to the equivalent view of graphs with distinguished vertices, we wish to compute $|\text{Hom}((G, \bar{x}) \rightarrow (H, \bar{y}))|$, modulo 2.

By definition of the enumeration (up to isomorphism) $\bar{y}_1, \dots, \bar{y}_\lambda$, there is some p such that $(H, \bar{y}) \cong (H, \bar{y}_p)$. Let \mathbf{v} be the vector that has a 1 in position p and has 0 in every other position. By Corollary 3.12, \mathbf{v} is implemented by some sequence $(\Theta_1, \bar{x}_1), \dots, (\Theta_t, \bar{x}_t)$ of graphs with r -tuples of distinguished vertices.

For each $i \in [t]$, let (G_i, \bar{x}) be the graph that results from taking the union of disjoint copies of G and Θ_i and identifying the j th element of \bar{x} with the j th element of \bar{x}_i for each $j \in [t]$. We have

$$\begin{aligned} \mathbf{v}_H(G, \bar{x}) \otimes \mathbf{v} &= \mathbf{v}_H(G, \bar{x}) \otimes \mathbf{v}_H((\Theta_1, \bar{x}_1) + \dots + (\Theta_t, \bar{x}_t)) \\ &= \bigoplus_{i \in [t]} (\mathbf{v}_H(G, \bar{x}) \otimes \mathbf{v}_H(\Theta_i, \bar{x}_i)) \\ &= \bigoplus_{i \in [t]} \mathbf{v}_H(G_i, \bar{x}). \end{aligned}$$

Now, sum the components of the vectors on the two sides of the equation. On the right, by Lemma 3.13, we have a value congruent modulo 2 to $\sum_{i \in [t]} |\text{Hom}(G_i \rightarrow H)|$. This can be computed by making t calls to an oracle for $\oplus \text{HOMSTO}H$, and t is bounded above by a constant, since H is fixed. On the left, we have, $|\text{Hom}((G, \bar{x}) \rightarrow (H, \bar{y}))|$, modulo 2, which is what we wish to compute. \square

The result we have proved appears similar to [Göbel et al. 2014, Theorem 3.2] but there is an important difference. In [Göbel et al. 2014], we wished to pin r vertices of G , each to the orbit of a vertex of H . In this paper, we focus on the problem $\oplus \text{PARTLABHOMSTO}H$, where we pin vertices of G to individual vertices of H . In order to achieve this, we essentially pin an r -tuple of vertices of G to the orbit of an r -tuple of vertices in H . To see the difference, consider the graph H in Figure 2. The orbits of single vertices are $\{a_1, a_2, a_3\}, \dots, \{d_1, d_2, d_3\}$. There are six homomorphisms from the single edge (x, y) to H that map x to the orbit of a_1 and y to the orbit of d_1 but only three that map the pair (x, y) to the orbit of the pair (a_1, d_1) , which is $\{(a_1, d_1), (a_2, d_2), (a_3, d_3)\}$.

4. HARDNESS GADGETS

In this section, we define gadgets that we will use to prove $\oplus\text{P}$ -completeness of $\oplus\text{HOMSTO}H$ problems, by reduction from the parity independent set problem $\oplus\text{IS}$, i.e., the problem of computing the number of independent sets in an input graph, modulo 2. $\oplus\text{IS}$ was shown to be $\oplus\text{P}$ -complete by Valiant [2006].

The gadgets we use are considerably more general than the ones we defined for cactus graphs in [Göbel et al. 2014]. This allows us to quickly prove hardness for large classes of square-free graphs and even to find gadgets non-constructively.

In fact, our definition of hardness gadgets and the proof that $\oplus\text{HOMSTO}H$ is $\oplus\text{P}$ -complete if H is involution-free and has a hardness gadget (Section 4.1) does not require the graphs to be square-free. However, whenever we find a gadget for a particular graph, it involves the “caterpillar gadgets” we introduce in Section 4.2. These gadgets do depend on H being square-free, as we show in Section 4.3.

4.1. $\oplus\text{P}$ -completeness

We now define the gadgets we use to prove hardness and show that they serve this purpose. Recall that a partially H -labelled graph J consists of an underlying graph $G(J)$ and a pinning function $\tau(J)$. In the discussion that follows, we will choose a set $\Omega_y \subseteq V(H)$ and a vertex $i \in \Omega_y$. Given a graph G whose independent sets we wish to count modulo 2, we will construct a partially H -labelled graph J and consider homomorphisms from J to H . $G(J)$ will contain a copy of $V(G)$ and we will be interested in homomorphisms that map every vertex in this copy to Ω_y . Vertices mapped to i will be in the independent set under consideration; vertices mapped to $\Omega_y - i$ will not be in the independent set.

Definition 4.1. A hardness gadget $(i, s, (J_1, y), (J_2, z), (J_3, y, z))$ for a graph H consists of vertices i and s of 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

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

The properties that we require are the following.

- (1) $|\Omega_y|$ is even and $i \in \Omega_y$.
- (2) $|\Omega_z|$ is even and $s \in \Omega_z$.
- (3) For each $o \in \Omega_y - i$ and each $x \in \Omega_z - s$, $|\Sigma_{o,x}|$ is even.
- (4) $|\Sigma_{i,s}|$ is odd and, for each $o \in \Omega_y - i$ and each $x \in \Omega_z - s$, $|\Sigma_{o,s}|$ and $|\Sigma_{i,x}|$ are odd.

Before proving that hardness gadgets give $\oplus\text{P}$ -completeness, we introduce some notation. Given partially H -labelled graphs $J_1 = (G_1, \tau_1)$ and $J_2 = (G_2, \tau_2)$, with $\text{dom}(\tau_1) \cap \text{dom}(\tau_2) = \emptyset$, we write $J_1 \cup J_2$ for the partially labelled graph $J' = (G', \tau')$, where $G' = G_1 \cup G_2$ and $\tau' = \tau_1 \cup \tau_2$. That is, $\text{dom}(\tau') = \text{dom}(\tau_1) \cup \text{dom}(\tau_2)$ and

$$\tau'(v) = \begin{cases} \tau_1(v) & \text{if } v \in \text{dom}(\tau_1) \\ \tau_2(v) & \text{if } v \in \text{dom}(\tau_2). \end{cases}$$

We will use the following notation to build partially labelled graphs containing many copies of some subgraph. For any “tag” T (which we will treat just as an arbitrary string) and any partially labelled graph J , denote by J^T a copy of J with every vertex $v \in V(G(J))$ renamed v^T .




Fig. 3. The construction of the partially labelled graphs K and J from an example graph G , as in the proof of Theorem 4.2.

THEOREM 4.2. *If an involution-free graph H has a hardness gadget then $\oplus\text{HOMSTO}H$ is $\oplus\text{P}$ -complete.*

PROOF. Let $(i, s, (J_1, y), (J_2, z), (J_3, y, z))$ be the hardness gadget for H and recall the sets Ω_y and Ω_z from Definition 4.1. We show how to reduce $\oplus\text{IS}$ to $\oplus\text{PARTLABHOMSTO}H$; the result then follows from Theorem 3.1 and $\oplus\text{P}$ -completeness of $\oplus\text{IS}$ [Valiant 2006]. Given an input graph G to $\oplus\text{IS}$, we construct an appropriate partially H -labelled graph J and show that $|\mathcal{I}(G)| \equiv |\text{Hom}(J \rightarrow H)| \pmod{2}$, where $\mathcal{I}(G)$ is the set of independent sets in G .

We construct J in two stages (see Figure 3). Take the union of disjoint copies $J_3^{e,v}$ of J_3 for every edge $e \in G$ and each endpoint v of e . For each edge $e = (u, v) \in G$, identify the vertices $z^{e,u}$ and $z^{e,v}$ and call this z^e . For each vertex $v \in G$, identify all the vertices $y^{e,v}$ such that e has v as an endpoint, and call this y^v . Call the resulting graph K .

To make J , take K and add a disjoint copy J_1^v of J_1 for every vertex $v \in G$ and a disjoint copy J_2^e of J_2 for every edge $e \in G$. For each vertex $v \in G$, identify the vertex y^v in K with the vertex y^v in J_1^v . For each edge $e = (u, v)$ in G , identify the vertex z^e in K with the vertex z^e in J_2^e .

We now proceed to show that $|\text{Hom}(J \rightarrow H)| \equiv |\mathcal{I}(G)| \pmod{2}$.

For a homomorphism $\sigma \in \text{Hom}(K \rightarrow H)$, let $[\sigma]$ be the set of extensions of σ to homomorphisms from J to H , i.e.,

$$[\sigma] = \{\sigma' \in \text{Hom}(J \rightarrow H) \mid \sigma(v) = \sigma'(v) \text{ for all } v \in V(G(K))\}.$$

Every homomorphism from J to H is the extension of a unique homomorphism from K to H , so we have

$$|\text{Hom}(J \rightarrow H)| = \sum_{\sigma \in \text{Hom}(K \rightarrow H)} |[\sigma]|. \quad (3)$$

From the structure of J , we have

$$|[\sigma]| = \left(\prod_{v \in V(G)} |\text{Hom}((J_1, y) \rightarrow (H, \sigma(y^v)))| \right) \left(\prod_{e \in E(G)} |\text{Hom}((J_2, z) \rightarrow (H, \sigma(z^e)))| \right).$$

By Definition 4.1, $|\text{Hom}((J_1, y) \rightarrow (H, a))|$ is odd if and only if $a \in \Omega_y$ and $|\text{Hom}((J_2, z) \rightarrow (H, b))|$ is odd if and only if $b \in \Omega_z$. Therefore, $|\sigma|$ is odd if and only if σ maps every vertex y^v into Ω_y and every z^e into Ω_z : call such a homomorphism “legitimate” (with respect to J_1 and J_2). We can rewrite (3) as

$$|\text{Hom}(J \rightarrow H)| \equiv |\{\sigma \in \text{Hom}(K \rightarrow H) \mid \sigma \text{ is legitimate}\}| \pmod{2}, \quad (4)$$

and, from this point, we restrict our attention to legitimate homomorphisms.

Given a legitimate homomorphism $\sigma \in \text{Hom}(K \rightarrow H)$, let $\sigma|_Y$ be the restriction of σ to the domain $\{y^v \mid v \in V(G)\}$. Write $\sigma \sim_Y \sigma'$ if $\sigma|_Y = \sigma'|_Y$ and write $[\sigma]_Y$ for the \sim_Y -equivalence class of σ . The classes $[\sigma]_Y$ partition the legitimate homomorphisms from K to H . We have

$$|[\sigma]_Y| = \prod_{(u,v) \in E(G)} n(\sigma(u), \sigma(v)),$$

where

$$n(a, a') = \sum_{b \in \Omega_z} |\text{Hom}((J_3, y, z) \rightarrow (H, a, b))| |\text{Hom}((J_3, y, z) \rightarrow (H, a', b))|.$$

By Definition 4.1, $|\Omega_z|$ is even, so the sum defining $n(a, a')$ has an even number of terms. $|\text{Hom}((J_3, y, z) \rightarrow (H, a, b))| = |\Sigma_{a,b}|$ is even if $a \in \Omega_y - i$ and $b \in \Omega_z - s$ and odd, otherwise. If $a = a' = i$, every term is odd and $n(a, a')$ is even; otherwise, exactly one term ($b = s$) is odd, so $n(a, a')$ is odd. Therefore, $|\sigma|$ is odd if and only if σ does not map a pair of adjacent vertices to i : that is, if the set $I(\sigma) = \{v \in V(G) \mid \sigma(y^v) = i\}$ is an independent set in G .

Choose representatives $\sigma_1, \dots, \sigma_k$, one from each \sim_Y -equivalence class. We have

$$\begin{aligned} |\text{Hom}(J \rightarrow H)| &\equiv |\{\sigma \in \text{Hom}(K \rightarrow H) \mid \sigma \text{ is legitimate}\}| \pmod{2} \\ &= \sum_{j=1}^k |[\sigma_j]_Y| \\ &\equiv |\{j \in [k] \mid I(\sigma_j) \text{ is independent}\}| \pmod{2} \\ &= \sum_{X \in \mathcal{I}(G)} |\{\sigma_j \mid j \in [k] \text{ and } I(\sigma_j) = X\}| \\ &\equiv |\mathcal{I}(G)| \pmod{2}, \end{aligned}$$

where the final equivalence is because the number of σ_j such that $I(\sigma) = X$ is exactly $|\Omega_y - i|^{|V(G) \setminus X|}$, which is odd because $|\Omega_y|$ is even. \square

4.2. Caterpillar gadgets

All our hardness gadgets use the following “caterpillar gadgets” as J_3 . We will also use two further kinds of gadget, “neighbourhood gadgets” and “ ℓ -cycle gadgets”, but we defer their definitions to the sections where they are used. As we will see in the following section, caterpillar gadgets rely on H being square-free.

Definition 4.3. For a path $P = v_0 \dots v_k$ in H of length at least 1, define the *caterpillar gadget* $J_P = (G, \tau)$ as follows (see Figure 4). $V(G) = \{u_1, \dots, u_{k-1}, w_1, \dots, w_{k-1}, y, z\}$




Fig. 4. The caterpillar gadget corresponding to a path $v_0 \dots v_k$. The vertices w_1, \dots, w_{k-1} in the gadget are pinned to vertices v_1, \dots, v_{k-1} in H , respectively.

and G is the path $yu_1 \dots u_{k-1}z$ together with edges (u_j, w_j) for $1 \leq j \leq k-1$. $\tau = \{w_1 \mapsto v_1, \dots, w_{k-1} \mapsto v_{k-1}\}$.

Note that, if P is a single edge, $G(J_P)$ is also the single edge (y, z) and $\tau(J_P) = \emptyset$.

In the following, we will repeatedly make use of the following fact about square-free graphs: if two distinct vertices have a common neighbour, they must have a unique common neighbour, since a pair of vertices with two common neighbours would form a 4-cycle.

LEMMA 4.4. *Let H be a square-free graph, let $k > 0$ and let $P = v_0 \dots v_k$ be a path in H .*

- (1) *For any $a \in \Gamma_H(v_0) - v_1$ and $\sigma \in \text{Hom}((J_P, y) \rightarrow (H, a))$, $\sigma(u_j) = v_{j-1}$ for all $j \in [k-1]$.*
- (2) *For any $b \in \Gamma_H(v_k) - v_{k-1}$ and $\sigma \in \text{Hom}((J_P, z) \rightarrow (H, b))$, $\sigma(u_j) = v_{j+1}$ for all $j \in [k-1]$.*

PROOF. The result is trivial for $k = 1$ so we assume $k > 1$. We prove the first part, by induction on j . The second part follows by symmetry (call the vertices on the path $v_k \dots v_0$ instead of $v_0 \dots v_k$).

First, take $j = 1$. From the structure of J_P , $\sigma(u_1)$ must be a neighbour of $\sigma(y) = a$ and of v_1 , which are distinct vertices. v_0 is a common neighbour of a and v_1 , so it must be their unique common neighbour, so $\sigma(u_1) = v_0$. Now, suppose that $\sigma(u_{j-1}) = v_{j-2}$. As in the base case, $\sigma(u_j)$ must be some neighbour of v_{j-2} and v_j , which are distinct. v_{j-1} is such a vertex, so it is the unique such vertex. \square

LEMMA 4.5. *Let H be a square-free graph. Let $k > 0$ and let $P = v_0 \dots v_k$ be a path in H with $\deg_H(v_j)$ odd for all $j \in \{1, \dots, k-1\}$. Let $\Omega_y \subseteq \Gamma_H(v_0)$ and $\Omega_z \subseteq \Gamma_H(v_k)$, with $i = v_1 \in \Omega_y$ and $s = v_{k-1} \in \Omega_z$. For each $o \in \Omega_y - i$ and each $x \in \Omega_z - s$:*

- (1) $|\text{Hom}((J_P, y, z) \rightarrow (H, o, x))| = 0$,
- (2) $|\text{Hom}((J_P, y, z) \rightarrow (H, o, s))| = 1$,
- (3) $|\text{Hom}((J_P, y, z) \rightarrow (H, i, x))| = 1$ and
- (4) $|\text{Hom}((J_P, y, z) \rightarrow (H, i, s))|$ is odd.

PROOF. If $k = 1$, $i = v_1$, $s = v_0$, $G(J_P)$ is the single edge (y, z) and $\tau(J_P) = \emptyset$. For any $o \in \Omega_y - i$ and $x \in \Omega_z - s$, we have $(o, s), (i, s), (i, x) \in E(H)$ so $(o, x) \notin E(H)$ because H is square-free. Parts 1–4 are immediate. For the remainder of the proof, we may assume that $k \geq 2$. Note that when $k = 2$, $i = s = v_1$ and this is the unique common neighbour of v_0 and v_2 in H .

For part 1, suppose, towards a contradiction, that $\sigma \in \text{Hom}((J_P, y, z) \rightarrow (H, o, x))$. In particular, $\sigma \in \text{Hom}((J_P, y) \rightarrow (H, o))$ so, by Lemma 4.4(1), $\sigma(u_1) = v_0$. We also have

$\sigma \in \text{Hom}((J_P, z) \rightarrow (H, x))$ so, by Lemma 4.4(2), $\sigma(u_1) = v_2$. But P is a simple path so $v_0 \neq v_2$.

For part 2, let $\sigma \in \text{Hom}((J_P, y, z) \rightarrow (H, o, s))$. Since $\sigma \in \text{Hom}((J_P, y) \rightarrow (H, o))$, $\sigma(u_j) = v_{j-1}$ for all $i \in [k-1]$ by Lemma 4.4(1). But now, σ is completely determined, so it is the unique element of $\text{Hom}((J_P, y, z) \rightarrow (H, o, s))$. Part 3 follows similarly from Lemma 4.4(2).

For part 4, first note that there is a homomorphism $\sigma^+ \in \text{Hom}((J_P, y, z) \rightarrow (H, i, s))$ with $\sigma^+(u_j) = v_{j+1}$ for all $j \in [k-1]$. Now, for $m \in [k-1]$, let

$$S_m = \{\sigma \in \text{Hom}((J_P, y, z) \rightarrow (H, i, s)) \mid m \text{ is minimal such that } \sigma(u_m) \neq v_{m+1}\}.$$

The sets $\{\sigma^+\}$ and S_1, \dots, S_{k-1} partition $\text{Hom}((J_P, y, z) \rightarrow (H, i, s))$.

We claim that, for any $\sigma \in S_m$, $\sigma(u_j) = v_{j-1}$ for all $j > m$. This is trivial for S_{k-1} so let $\sigma \in S_m$ with $m < k-1$. $\sigma(u_{m+1})$ must be a neighbour of both $\sigma(u_{m+2}) = v_{m+2}$ and $\sigma(u_m) \in \Gamma_H(v_m)$. By definition of S_m , these are distinct vertices so v_m is their unique common neighbour and so $\sigma(u_{m+1}) = v_m$. Now, if $\sigma(u_j) = v_{j-1}$ for some $j \in \{m+1, \dots, k-2\}$, then $\sigma(u_{j+1})$ must be a neighbour of both $\sigma(u_{j+2}) = v_{j+2}$ and v_{j-1} : v_j is the unique such vertex, so $\sigma(u_{j+1}) = v_j$. This establishes the claim.

But, now, for any $\sigma \in S_m$, we have $\sigma(u_j) = v_{j+1}$ for $j < m$ and $\sigma(u_j) = v_{j-1}$ for $j > m$. $\sigma(y) = i$, $\sigma(z) = s$ and $\sigma(w_j) = v_j$ for each $j \in [k-1]$. Finally, $\sigma(u_m)$ may take any value in $\Gamma_H(v_m) - v_{m+1}$. It follows that, for all m , $|S_m| = \deg_H(v_m) - 1$, which is even. $|\text{Hom}((J_P, y, z) \rightarrow (H, i, s))| = 1 + \sum_m |S_m|$, which is odd, as required. \square

4.3. Caterpillar gadgets and 4-cycles

Before proceeding to find hardness gadgets for square-free graphs in the next section, we pause to show why 4-cycles cause problems for caterpillar gadgets and, in particular, why Lemma 4.5 does not apply to graphs containing 4-cycles.

Consider first the one-edge caterpillar gadget J_1 associated with the path v_0v_1 in the graph H_1 in Figure 5. This corresponds to $k = 1$ in Lemma 4.5 and we have $i = v_1$ and $s = v_0$. Taking $\Omega_y = \Gamma_{H_1}(v_0) = \{v'_0, v_1\}$ and $\Omega_z = \Gamma_{H_1}(v_1) = \{v_0, v'_1\}$ satisfies the conditions of the lemma. However, taking $o = v'_0 \in \Omega_y - i$ and $x = v'_1 \in \Omega_z - s$, we have $|\text{Hom}((J_1, y, z) \rightarrow (H, o, x))| = 1$ so part 1 of the lemma does not hold. However, the other three parts hold, as

$$\begin{aligned} |\text{Hom}((J_1, y, z) \rightarrow (H, o, s))| &= |\text{Hom}((J_1, y, z) \rightarrow (H, i, x))| \\ &= |\text{Hom}((J_1, y, z) \rightarrow (H, i, s))| = 1. \end{aligned}$$

Now, consider longer paths such as the path $P = v_0 \dots v_k$ in H_k in Figure 5, for some $k \geq 2$. The associated caterpillar gadget J_P is also shown in the figure. For each $j \in \{1, \dots, k-1\}$, $\deg_{H_k}(v_i)$ is odd. We have $i = v_1$ and $s = v_{k-1}$ (with $i = s$ in the case $k = 2$). Again, take $\Omega_y = \Gamma_{H_k}(v_0) = \{v'_0, v_1\}$, take $\Omega_z = \Gamma_{H_k}(v_k) = \{v_{k-1}, v'_k\}$ and take $o = v'_0 \in \Omega_y - i$ and $x = v'_k \in \Omega_z - s$.

Once again part 1 of the lemma fails. We have $|\text{Hom}((J_P, y, z) \rightarrow (H_k, o, x))| = 1$, since there is a homomorphism that maps u_j to v'_j for each $j \in \{1, \dots, k-1\}$. This is the only possible homomorphism from (J_P, y, z) to (H_k, o, x) since there is only one k -path from o to x that the k -path in J_P can be mapped to. For a hardness gadget, it would suffice for $|\text{Hom}((J_P, y, z) \rightarrow (H_k, o, x))|$ to be even (not necessarily zero) but it is odd for every k .

For H_k , the other parts of the lemma fail, too. We have

$$|\text{Hom}((J_P, y, z) \rightarrow (H, o, s))| = |\text{Hom}((J_P, y, z) \rightarrow (H, i, x))| = k.$$

When the target is (H, o, s) , the k -path in J_P can be mapped to any of the k k -paths in H_k from o to s (following along $v'_0v'_1 \dots$ and then dropping down along an edge v'_jv_j




Fig. 5. Examples of graphs containing 4-cycles for which caterpillar gadgets (Definition 4.3 and Lemma 4.5) fail. The graphs H_1 and H_k ($k \geq 2$) are shown, along with the caterpillar gadgets J_1 and J_P , corresponding to the paths v_0v_1 and $v_0 \dots v_k$, respectively. The labels o , s , i and x are referenced in the text.


and then following $v_jv_{j+1} \dots v_{k-1}$). The case with target (H, i, x) is similar. So in both cases, the number of homomorphisms is k . When k is odd, this is not a real problem. The purpose of Lemma 4.5 is to show that caterpillar gadgets can be used as J_3 in a hardness gadget, and the definition of hardness gadgets only requires that $|\Sigma_{o,s}|$ and $|\Sigma_{i,x}|$ (i.e., $|\text{Hom}((J_P, y, z) \rightarrow (H, o, s))|$ and $|\text{Hom}((J_P, y, z) \rightarrow (H, i, x))|$, respectively) be odd and not necessarily 1. However, this relaxation doesn't help when k is even.

Finally, for part 4, consider a homomorphism from (J_P, y, z) to (H, i, s) . The image of the path $yu_1 \dots u_{k-1}z$ in H must be a k -walk $v_1x_1 \dots x_{k-1}v_{k-1}$ with the property that x_j is adjacent to v_j for each $j \in \{1, \dots, k-1\}$. This means that $x_j \in \{v_{j-1}, v'_j, v_{j+1}\}$. There are two kinds of k -walk satisfying these criteria. The first kind uses only the vertices $\{v_0, \dots, v_k\}$. Such a walk must be either $v_1v_0v_1v_2 \dots v_{k-1}$ or $v_1 \dots v_\alpha v_{\alpha+1} v_\alpha \dots v_{k-1}$ for some $\alpha \in \{1, \dots, k-1\}$. The second kind uses some of the vertices $\{v'_1, \dots, v'_{k-1}\}$. Such a walk must be of the form $v_1 \dots v_\alpha v'_\alpha \dots v'_\beta v_\beta \dots v_{k-1}$ for some $1 \leq \alpha \leq \beta \leq k-1$. There are k walks of the first kind and $\frac{1}{2}k(k-1)$ of the second. Thus,

$$|\text{Hom}((J_1, y, z) \rightarrow (H, i, s))| = k + \frac{1}{2}k(k-1) = \frac{1}{2}k(k+1),$$

which is odd if and only if k is congruent to 1 or 2, mod 4 but is required to be odd for all k .

We note that $\oplus \text{HOMSTO}H_1$ is $\oplus \text{P}$ -complete, as is $\oplus \text{HOMSTO}H_k$, for every $k \geq 2$. H_1 is an involution-free cactus graph with more than one vertex so it is hard by the main theorem of [Göbel et al. 2014]. We claim that $\mathcal{X} = (i, s, (J_1, y), (J_2, z), (J_3, y, z))$, as shown in Figure 6, is a hardness gadget for H_k . We have $\Omega_y = \{v_0, v'_1\} = \{o, i\}$ and $\Omega_z = \{v_1, v'_2\} = \{s, x\}$: both are even and $i \in \Omega_y$ and $s \in \Omega_z$. There is no edge ox in H_k so $|\Sigma_{o,x}| = 0$, which is even. There are edges os , ix and is in H_k , so $|\Sigma_{o,s}| = |\Sigma_{i,x}| = |\Sigma_{i,s}| = 1$, which is odd. This establishes that \mathcal{X} is a hardness gadget so, since H_k is involution-free, $\oplus \text{HOMSTO}H_k$ is $\oplus \text{P}$ -complete by Theorem 4.2. Ironically, the part J_3 of \mathcal{X} is the one-edge caterpillar gadget associated with the path $v_1v'_1$ in H_k . The failure of Lemma 4.5 in the presence of 4-cycles only means that caterpillar gadgets are not guaranteed to work, not that they never work.

Fig. 6. A hardness gadget for the graph H_k (see also Figure 5).

5. FINDING HARDNESS GADGETS

In this section, we show how to find hardness gadgets for all connected, involution-free, square-free graphs. The simplest case is when the graph contains at least two vertices of even degree. Faben and Jerrum [2015] used the fact that all involution-free trees have at least two even-degree vertices, though we use different gadgets because we are dealing with graphs containing cycles as well as trees. For graphs with only one even-degree vertex, we show that an appropriate vertex deletion produces a component with more than one even-degree vertex and show how to simulate such a vertex deletion using gadgets.

This leaves graphs where every vertex has odd degree. In Section 5.2, we show how to use odd-length cycles to find a hardness gadget. The remaining case, bipartite graphs in which every vertex has odd degree, is covered in Section 5.3, where we use Corollary 3.7, our version of Lovász’s result, to non-constructively demonstrate that a hardness gadget always exists.

We will use the following fact.

LEMMA 5.1. *An involution-free graph with at least two vertices but at most one even-degree vertex contains a cycle.*

PROOF. We prove the contrapositive. Let G be an involution-free acyclic graph. At most one component of G is an isolated vertex so, if G has two or more vertices, it has at least one component with two or more vertices. This component is an involution-free tree which, by [Faben and Jerrum 2015, Lemma 5.3], contains at least two vertices of even degree. \square

5.1. Even-degree vertices

We prove that involution-free graphs containing at least one vertex of positive, even degree have a hardness gadget. In this section, we will use one extra kind of gadget.

Definition 5.2. For a vertex $v \in V(H)$, define the *neighbourhood gadget* $J_{\Gamma(v),x} = (G, \{w \mapsto v\})$, where G is the single edge (x, w) .

It is immediate from the definition that, for any $v \in V(H)$,

$$|\text{Hom}((J_{\Gamma(v),x}, x) \rightarrow (H, u))| = \begin{cases} 1 & \text{if } u \in \Gamma_H(v) \\ 0 & \text{otherwise.} \end{cases}$$

We first show how to find hardness gadgets for connected graphs containing at least two even-degree vertices (their degree must be positive, since the graph is connected) and then deal with the harder case of graphs containing exactly one vertex of posi-

tive, even degree. The following lemma constructs a caterpillar gadget, so the lemma depends on H being square-free. The extended conclusion about pinned vertices is needed for technical reasons in the proof of Lemma 5.7.

LEMMA 5.3. *Let H be a connected, square-free graph with at least two even-degree vertices. Then H has a hardness gadget $(i, s, (J_1, y), (J_2, z), (J_3, y, z))$. Furthermore, we can choose J_1, J_2 and J_3 so that each contains at least one pinned vertex.*

PROOF. Let $v_0 \dots v_m$ be a path in H between distinct even-degree vertices v_0 and v_m and let $P = v_0 \dots v_k$, where $k \in \{1, \dots, m\}$ is minimal such that $\deg_H(v_k)$ is even. We claim that $(v_1, v_{k-1}, (J_{\Gamma(v_0), y}, y), (J_{\Gamma(v_k), z}, z), (J_P, y, z))$ is a hardness gadget. $|\Omega_y|$ and $|\Omega_z|$ are even because v_0 and v_k have even degree; and they contain v_1 and v_{k-1} , respectively. The remaining properties required by Definition 4.1 hold by Lemma 4.5, since v_1, \dots, v_{k-1} have odd degree.

Each of $J_{\Gamma(v_0), y}$ and $J_{\Gamma(v_k), z}$ contains a pinned vertex and, if $k > 1$, J_P also contains at least one pinned vertex. If $k = 1$, then $G(J_P)$ is the single edge (y, z) and $\tau(J_P) = \emptyset$. However, we may add to $G(J_P)$ a new vertex w_0 and an edge (w_0, y) and set $\tau(J_P) = \{w_0 \mapsto v_0\}$: this requires y to be mapped to a neighbour of v_0 . This has no effect on the hardness gadget since Definition 4.1 only imposes requirements on $|\text{Hom}((J_3, y, z) \rightarrow (H, a, b))|$ when $a \in \Omega_y$. Since $\Omega_y = \Gamma_H(v_0)$, we are already only considering homomorphisms that map y to a neighbour of v_0 and the change to J_3 is merely restating this condition. \square

It is worth noting that, since all involution-free trees have at least two even-degree vertices, Lemma 5.3 implies Faben and Jerrum's dichotomy for $\oplus\text{HOMSTO}H$ where H is a tree [2015]. They also use two even-degree vertices but their gadgets rely on the fact that there is a unique path between two vertices of a tree, which doesn't hold in general graphs. However, from Lemma 5.3, we conclude that uniqueness of the path is not required and we can prove hardness even when there are multiple paths between even-degree vertices.

To handle graphs with fewer than two vertices of even degree, we first investigate the results of deleting vertices from such graphs. If we delete the unique even-degree vertex from a connected graph, then each component of the resulting graph contains at least one vertex of even degree. If we are lucky, one of the resulting components will contain two or more vertices of even degree, raising the hope that we can use Lemma 5.3 to prove $\oplus\text{P}$ -completeness. If all of the resulting components have exactly one even-degree vertex, then we can iterate, deleting those vertices to obtain yet more fragments. As long as the graph contains at least one cycle, it is not hard to see that we can eventually obtain a component with two or more even-degree vertices. However, to apply Lemma 5.3, we must ensure that the resulting component has no involution. We prove this in the following two lemmas.

LEMMA 5.4. *Let H be an involution-free graph with exactly one vertex v of positive, even degree. Then $H' = H - v$ is also involution-free.*

PROOF. Each vertex $u \in \Gamma_H(v)$ has odd degree in H and has exactly one neighbour removed, so $\deg_{H'}(u)$ is even. Suppose, towards a contradiction, that ρ is an involution of H' . No automorphism can map an odd-degree vertex to an even-degree vertex or vice-versa and $\Gamma_H(v)$ is exactly the set of even-degree vertices in H' . Therefore, the restriction of ρ to the neighbours of v is a permutation. Define $\hat{\rho}: V(H) \rightarrow V(H)$ by $\hat{\rho}(v) = v$ and $\hat{\rho}(w) = \rho(w)$ for $w \neq v$. $\hat{\rho}$ preserves all edges in H' and all edges incident on v in H , so it is an involution of H , contradicting the supposition that H has no involution. \square

So far, we have described our goal as being to iteratively delete vertices until we find a component with more than one even-degree vertex. This is a useful intuition but we do not know how to simulate such a sequence of vertex deletions using gadgets. Instead, we show how to achieve the goal of a component with more than one even-degree vertex by deleting a set of vertices, which we do know how to do with a gadget.

For a vertex $v \in V(H)$ and an integer $r \geq 0$, let $B_r(v) = \{u \in V(H) \mid \text{dist}(u, v) = r\}$.

COROLLARY 5.5. *Let H be an involution-free graph that has exactly one vertex v of positive, even degree. For some r , $H - B_r(v)$ has an involution-free component H^* that does not contain v but does contain at least two even-degree vertices. Furthermore, we can take $r = \min \{\text{dist}(v, w) \mid w \text{ is on a cycle}\}$.*

PROOF. H contains a cycle by Lemma 5.1 so we can take r as in the statement of the lemma and this is well-defined. If $r = 0$, then v is in some cycle C in H . $H - v$ has no involution by Lemma 5.4, so no component of $H - v$ has an involution. The component H^* of $H - v$ that contains $C - v$ contains at least two vertices of $\Gamma_H(v)$ (v 's two neighbours in C) and these vertices have even degree in H^* . H^* does not, of course, contain v .

Suppose that $r > 0$. By the choice of r , there must be a component H' of $H - B_{r-1}(v)$ that contains a vertex $v_r \in B_r(v)$ that is in a cycle C' of H' . Since no vertex at distance less than r from v is in a cycle in H , there is a unique path from v to v_r . Let this be $v_0 \dots v_r$, where $v = v_0$. A simple induction on $j = 0, \dots, r-1$, using Lemma 5.4, shows that the component of $H - v_j$ containing v_r has no involution, does not contain v and has exactly one even-degree vertex: namely, v_{j+1} . In particular, the component of $H - v_{r-1}$ that contains v_r is H' . But, now, the component of $H' - v_r$ that contains $C' - v_r$ has no involution (because no component of $H' - v_r$ has an involution) and contains at least two vertices of even degree (because v_r has at least two neighbours in C'). Further, this component is the component H^* of $H - B_r(v)$ that we seek. \square

Thus, starting with an involution-free graph H containing only one vertex of positive, even degree, we have shown how to make a set of vertex deletions (some set $B_r(v)$) to obtain an involution-free component H^* with at least two even-degree vertices. We now show that we can achieve these vertex deletions using gadgetry. The following technical lemma allows us to construct a gadget that, in a sense, “selects” the vertices of H^* within H .

LEMMA 5.6. *Let H be a graph, let $P = x_0 \dots x_{r+1}$ with $r \geq 0$ be a path in H and let $w \in V(H)$. If every vertex in H within distance $r-1$ of w has odd degree, then $|\text{Hom}((P, x_0) \rightarrow (H, w))|$ has opposite parity to the number of distinct r -paths in H from w to vertices of even degree.*

PROOF. We prove the lemma by induction on r . For $r = 0$, the result is trivial. The condition on vertices within distance $r-1$ is vacuous. The number of 0-paths from w to vertices of even degree is zero if $\text{deg}(w)$ is odd; it is one if $\text{deg}(w)$ is even; and $|\text{Hom}((P, x_0) \rightarrow (H, w))| = \text{deg}(w)$.

Suppose the result holds for the path $P = x_0 \dots x_{r+1}$ and consider the path Px_{r+2} and a graph H in which every vertex within distance r of w has odd degree.

Every homomorphism σ from (Px_{r+2}, x_0) to (H, w) induces a homomorphism $\hat{\sigma}$ from (P, x_0) to (H, w) . Write $\sigma \sim \sigma'$ if $\hat{\sigma} = \hat{\sigma}'$. \sim is an equivalence relation and its equivalence classes partition $\text{Hom}((Px_{r+2}, x_0) \rightarrow (H, w))$. Let $[\sigma]$ be the \sim -equivalence class of σ .

If every vertex within distance r of w in H has odd degree, there are no r -paths from w to vertices of even degree so, by the inductive hypothesis, there are an odd number of homomorphisms from (P, x_0) to (H, w) , so there are an odd number of equivalence classes. Further, $|\text{Hom}((Px_{r+2}, x_0) \rightarrow (H, w))| = \text{deg}(\sigma(x_{r+1}))$ (this is well-defined since $\sigma(x_{r+1}) = \hat{\sigma}(x_{r+1})$), so

all homomorphisms $\sigma' \in [\sigma]$ agree on the value of $\sigma'(x_{r+1})$. Any vertex of even degree is at distance at least $r+1$ from $w = \sigma(x_0)$ so, if $\deg_H(\sigma(x_{r+1}))$ is even, then the r -walk $\sigma(x_0)\sigma(x_1)\dots\sigma(x_{r+1})$ is, in fact, a simple $(r+1)$ -path. Therefore, the number N of even-cardinality equivalence classes is equal to the number of $(r+1)$ -paths in H from w to a vertex of even degree, and subtracting these from the total number of equivalence classes gives $|\text{Hom}((Px_{r+2}, x_0) \rightarrow (H, w))| \equiv 1 - N \pmod{2}$, as required. \square

Now, we can obtain a hardness gadget for H by combining the “selection gadget” with the hardness gadget for the subgraph H^* given to us by Corollary 5.5.

LEMMA 5.7. *Any involution-free, square-free graph H that has exactly one vertex v of positive, even degree has a hardness gadget.*

PROOF. Let $r = \min\{\text{dist}(v, w) \mid w \text{ is on a cycle}\}$. By Corollary 5.5, there is an involution-free component H^* of $H - B_r(v)$ that does not contain v but contains at least two vertices of even degree. H^* is square-free because it is an induced subgraph of a square-free graph. Therefore, by Lemma 5.3, H^* has a hardness gadget $\mathcal{X}^* = (i, s, (J_1^*, y), (J_2^*, z), (J_3^*, y, z))$ in which each of J_1^* , J_2^* and J_3^* contains a pinned vertex.

We construct a hardness gadget \mathcal{X} for H from \mathcal{X}^* . Let P be a path of length $r+1 \geq 1$, with vertices $x_0 \dots x_{r+1}$. Let $J_1 = (G, \tau)$ be the partially H -labelled graph such that $\tau = \tau(J_1^*)$ and G is defined from $G(J_1^*)$ as follows: start with $G(J_1^*)$ and, for every vertex $u \in G(J_1^*)$, add a new copy of P and identify that copy’s vertex x_0 with u . Define J_2 and J_3 similarly, from J_2^* and J_3^* . We claim that the tuple

$$\mathcal{X} = (i, s, (J_1, y), (J_2, z), (J_3, y, z))$$

is the desired hardness gadget for H .

To find out what \mathcal{X} does, we first consider homomorphisms from one copy of the path P to H . For a vertex $w \in V(H)$, let $N_w = |\text{Hom}((P, x_0) \rightarrow (H, w))|$. If $\text{dist}(v, w) = r$ (i.e., $w \in B_r(v)$), then there is a unique r -path from w to a vertex of even degree. This is because v is the unique vertex of even degree and, if there were distinct r -paths Q_1 and Q_2 from w to v then $Q_1 \cup Q_2$ would contain a cycle, which would contain vertices at distance strictly less than r from v , contradicting the definition of r . If $\text{dist}(v, w) > r$, then there are no r -paths from w to even-degree vertices. Therefore, by Lemma 5.6, N_w is even if $\text{dist}(v, w) = r$ and N_w is odd if $\text{dist}(v, w) > r$ (we will see that the parity of N_w does not matter if $\text{dist}(v, w) < r$).

Now, let $a \in V(H)$ and consider homomorphisms $\sigma, \sigma' \in \text{Hom}((J_1, y) \rightarrow (H, a))$. Write $\sigma \sim \sigma'$ if $\sigma(u) = \sigma'(u)$ for all $u \in V(G(J_1^*))$ and write $[\sigma]$ for the \sim -equivalence class containing σ . $|\text{Hom}((J_1, y) \rightarrow (H, a))|$ is the sum of the sizes of the \sim -equivalence classes. For any σ , we have

$$|[\sigma]| = \prod_{x \in V(G(J_1^*))} |\text{Hom}((P, x_0) \rightarrow (H, \sigma(x)))|.$$

Therefore, $|\sigma|$ is even if σ maps any vertex of $G(J_1^*)$ into $B_r(v)$. In this case, $|\sigma|$ contributes nothing to the sum, modulo 2.

Thus, we may restrict our attention to homomorphisms from J_1^* to H that have no vertex in $B_r(v)$ in their image. J_1^* is connected and contains a vertex pinned to a vertex in H^* . Therefore, restricting to homomorphisms that have no vertex in $B_r(v)$ in their image means restricting to homomorphisms whose image is wholly within H^* . For any vertex $w \in H^*$, $\text{dist}_H(v, w) > r$, so this gives

$$|\text{Hom}((J_1, y) \rightarrow (H, a))| \equiv |\text{Hom}((J_1^*, y) \rightarrow (H^*, a))| \pmod{2},$$




Fig. 7. The ℓ -cycle gadget $J_{\ell,P,x}$ corresponding to a path $P = v_1 \dots v_k$ in an ℓ -cycle in H .

for any $a \in V(H^*)$ and $|\text{Hom}((J_1, y) \rightarrow (H, a))| \equiv 0 \pmod{2}$, for $a \notin V(H^*)$; and similarly for J_2 and J_3 . Thus, since \mathcal{X}^* is a hardness gadget for H^* , \mathcal{X} is a hardness gadget for H . \square

The proof of Lemma 5.7 does not explicitly use caterpillar gadgets. However, the hardness gadget \mathcal{X} is constructed from \mathcal{X}^* , which was produced by Lemma 5.3. It follows that J_3^* is a caterpillar gadget, so Lemma 5.7 requires H to be square-free, as stated.

5.2. Odd cycles

In the previous section, we showed how to find a hardness gadget for any involution-free, square-free graph containing at least one vertex of even degree. In this section, we show that any square-free graph in which all vertices have odd degree has a hardness gadget if it has an odd cycle. We first introduce a gadget for selecting certain vertices in cycles.

Definition 5.8. (See Figure 7). Let $P = v_1 \dots v_k$ be a path in H . For any $\ell > \max\{2, k\}$, define the ℓ -cycle gadget $J_{\ell,P,x} = (G, \tau)$ where G is the cycle $xu_1 \dots u_{\ell-1}x$ and $\tau = \{u_1 \mapsto v_1, \dots, u_k \mapsto v_k\}$.

Recall that the odd-girth of a graph is the length of its shortest odd cycle. By convention, the odd-girth of a graph without odd cycles is infinite; in the following, we write “a graph whose odd-girth is ℓ ” as a short-hand for “a graph whose odd-girth is finite and equal to ℓ .”

LEMMA 5.9. *Let H be a graph whose odd-girth is ℓ and let G be an ℓ -cycle. The image of G under any homomorphism from G to H is an ℓ -cycle in H .*

PROOF. Let $G = u_0 \dots u_{\ell-1}u_0$. Since G is an ℓ -cycle and H contains an ℓ -cycle, $\text{Hom}(G \rightarrow H)$ is non-empty so let $\sigma \in \text{Hom}(G \rightarrow H)$. Let C be the image of G under σ , i.e., subgraph of H consisting of vertices $\{\sigma(u_0), \dots, \sigma(u_{\ell-1})\}$ and edges $\{(\sigma(u_j), \sigma(u_{j+1})) \mid 0 \leq j < \ell\}$, with addition on indices carried out modulo ℓ . Suppose towards a contradiction that C is not an ℓ -cycle. Since C has at most ℓ vertices and at most ℓ edges, it cannot have an ℓ -cycle as a proper subgraph. Since H has no odd cycles shorter than ℓ , C must be bipartite. But then the walk $\sigma(u_0)\sigma(u_1) \dots \sigma(u_{\ell-1})\sigma(u_0)$ is an odd-length walk from a vertex to itself and no such walk can exist in a bipartite graph. \square

COROLLARY 5.10. *Let H be a graph whose odd-girth is ℓ . For any path P on fewer than ℓ vertices, $|\text{Hom}((J_{\ell,P,x}, x) \rightarrow (H, v))|$ is the number of ℓ -cycles in H that contain the path vP .*

PROOF. By Lemma 5.9, the image of $G(J_{\ell,P,x})$ under any homomorphism to H is an ℓ -cycle in H and, because of the pinning and distinguished vertex, this cycle must contain the path vP . \square

Let $\#C_{\ell}(vw)$ be the number of ℓ -cycles in H containing the edge (v, w) .

LEMMA 5.11. *Let H be a graph whose odd-girth is ℓ . Every vertex $v \in V(H)$ has an even number of neighbours w such that $\#C_{\ell}(vw)$ is odd.*

PROOF. If v is not in any ℓ -cycle, the claim is vacuous: the even number is zero. Otherwise, let $C = vw_1 \dots w_{\ell-1}v$ be an ℓ -cycle in H . If $w_j \in \Gamma_H(v)$ for some even $j \neq \ell - 1$, the odd cycle $vw_1 \dots w_j v$ contradicts the stated odd-girth of H . If $w_j \in \Gamma_H(v)$ for some odd $j \neq 1$, the odd cycle $vw_j \dots w_{\ell-1}v$ contradicts the odd-girth. Therefore, w_1 and $w_{\ell-1}$ are the only vertices in C that are adjacent to v and every ℓ -cycle through v contributes exactly 2 to $\sum_{w \in \Gamma_H(v)} \#C_{\ell}(vw)$. Therefore, the sum is even, so it has an even number of odd terms. \square

LEMMA 5.12. *Let H be a square-free graph whose odd-girth is ℓ . If H contains an edge that is in an odd number of ℓ -cycles, then H has a hardness gadget.*

Note that, for the case $\ell = 3$, any edge in a 3-cycle in H must be in exactly one 3-cycle since, if an edge (x, y) is in distinct 3-cycles $xyzx$ and $xyz'x$, then $xzyz'x$ is a 4-cycle in H , which is forbidden by the hypothesis of the lemma. The absence of 4-cycles is also required for the caterpillar gadget produced in the proof.

PROOF. Let (i, s) be an edge in an odd number of ℓ -cycles in H . Let J_1 be the ℓ -cycle gadget $J_{\ell,s,y}$ (so $\tau(J_1) = \{u_1 \mapsto s\}$) and let J_2 be the ℓ -cycle gadget $J_{\ell,i,z}$. Let $G(J_3)$ be the single edge (y, z) and let $\tau(J_3) = \emptyset$ (J_3 is, technically, a caterpillar gadget but it is easier to analyse it directly).

We claim that $(i, s, (J_1, y), (J_2, z), (J_3, y, z))$ is a hardness gadget for H . By Corollary 5.10, $|\text{Hom}((J_{\ell,s,y}, y) \rightarrow (H, v))|$ is the number of ℓ -cycles in H that contain the edge (v, s) , so

$$\Omega_y = \{v \in V(H) \mid (v, s) \text{ is in an odd number of } \ell\text{-cycles}\}.$$

Thus, $|\Omega_y|$ is even by Lemma 5.11. Ω_y contains i by the choice of the edge (i, s) in an odd number of ℓ -cycles. Similarly, Ω_z is even and contains s . To verify the remaining properties required by Definition 4.1, note that J_3 is a single edge so, for any $a, b \in V(H)$, $|\text{Hom}((J_3, y, z) \rightarrow (H, a, b))|$ is 1 if $(a, b) \in E(H)$ and 0, otherwise. We have $\Omega_y \subseteq \Gamma_H(s)$ and $\Omega_z \subseteq \Gamma_H(i)$ so, for any $o \in \Omega_y - i$ and any $x \in \Omega_z - s$, H contains the edges (o, s) , (s, i) and (i, x) but it cannot contain the edge (o, x) because H is square-free. \square

LEMMA 5.13. *Let H be a square-free graph in which every vertex has odd degree. If H contains an odd cycle, then it has a hardness gadget.*

PROOF. Let ℓ be the odd-girth of H . If H contains an edge in an odd number of ℓ -cycles (which is guaranteed for $\ell = 3$, since H is square-free), then H has a hardness gadget by Lemma 5.12. So, for the remainder of the proof, we may assume that the shortest odd cycle in H has length $\ell > 4$ and that every edge is in a (not necessarily positive) even number of ℓ -cycles.

Let $P = v_k v_{k+1} \dots v_{\ell-1} v_0$ be a longest path that is in a positive, even number of ℓ -cycles (see Figure 8; it turns out to be most convenient to label the vertices in this




Fig. 8. The parts J_1 , J_2 and J_3 of the hardness gadget constructed in the proof of Lemma 5.13. The corresponding cycle in H is indicated in grey within each gadget. The path $P = v_k \dots v_{l-1} v_0$ is undirected but the arrow indicates the order in which the vertices are listed.

order; the path has length $\ell - k$). Such a path certainly exists because any edge in an ℓ -cycle is in a positive, even number of them. So, in particular, P contains at least one edge. Further P has fewer than $\ell - 1$ edges, because any path on $\ell - 1$ edges is in at most one ℓ -cycle, since H has no parallel edges. Let $C = v_0 v_1 \dots v_{l-1} v_0$ be an ℓ -cycle containing P . Let $\text{rev}(P) = v_0 v_{l-1} \dots v_k$ be the path P with the vertices listed in the reverse order.

Let $i = v_1$ and $s = v_{k-1}$. Let J_1 be the ℓ -cycle gadget $J_{\ell, \text{rev}(P), y}$, let J_2 be the ℓ -cycle gadget $J_{\ell, P, z}$, and let J_3 be the caterpillar gadget $J_{v_0 \dots v_k}$.

We claim that $(i, s, (J_1, y), (J_2, z), (J_3, y, z))$ is a hardness gadget for H . Since P was chosen to be a longest path in a positive, even number of ℓ -cycles, any path uP in H must be in an odd number of ℓ -cycles or in none at all. Since P itself is in an even number of ℓ -cycles, the number of extensions uP in an odd number of cycles must be even. By Corollary 5.10, $|\text{Hom}((J_{\ell, P, z}, z) \rightarrow (H, u))|$ is the number of ℓ -cycles in H that contain the path uP . Therefore, $|\Omega_z|$ is precisely the set of vertices u such that uP is in an odd number of ℓ -cycles, so we have established that $|\Omega_z|$ is even. Since sP is an extension of P , it is not in a positive, even number of ℓ -cycles; it is in at least one ℓ -cycle (namely, C) so it is in an odd number of them. Therefore, $s \in \Omega_z$. Similarly, $|\Omega_y|$ is even and $i \in \Omega_y$.

It remains to verify that the conditions of Lemma 4.5 hold for J_3 , so that lemma gives us the remaining properties we need from Definition 4.1. All vertices in H have odd degree by assumption, including in particular the interior vertices of P . We have already established that $i = v_1 \in \Omega_y$ and $s = v_{k-1} \in \Omega_z$. Finally, $\Omega_y \subseteq \Gamma_H(v_0)$ because, in $G(J_1)$, y is adjacent to a vertex that is pinned to v_0 . Similarly, $\Omega_z \subseteq \Gamma_H(v_k)$. \square

5.3. Bipartite graphs

The only remaining case is bipartite graphs H in which every vertex has odd degree. We show that, if H has an “even gadget”, it has a hardness gadget. And it turns out that every connected bipartite graph with more than one edge has an even gadget.

Definition 5.14. An *even gadget* for a bipartite graph H with at least one edge is an edge (a, b) of H together with a connected bipartite graph G with a distinguished edge (w, x) such that $|\text{Hom}((G, w, x) \rightarrow (H, a, b))|$ is even.

Note that, for bipartite G and H , with edges (w, x) and (a, b) , respectively, there is always at least one homomorphism from (G, w, x) to (H, a, b) , since the whole of G can be mapped to the edge (a, b) . So, although Definition 5.14 only requires $|\text{Hom}((G, w, x) \rightarrow (H, a, b))|$ to be even, the number of homomorphisms is always non-zero.

Suppose that H is any connected bipartite graph with more than one edge such that, for some edge (a, b) of H , (H, a, b) is involution-free. We will show that H has an even gadget. If, furthermore, H is square-free, this even gadget gives a hardness gadget. If H is also involution-free, the hardness gadget implies $\oplus P$ -completeness of $\oplus \text{HOMSTO}H$, by Theorem 4.2.

LEMMA 5.15. *Suppose that H is a connected bipartite graph with more than one edge such that, for some edge (a, b) of H , (H, a, b) is involution-free. Then H has an even gadget.*

PROOF. Let H be a graph satisfying the conditions in the statement of the lemma. Let K_2 be the graph consisting of the single edge (a, b) . Clearly, (K_2, a, b) is involution-free (since there are no non-trivial automorphisms of K_2 that fix a and b) and $H \not\cong K_2$ since H has more than one edge, so $(H, a, b) \not\cong (K_2, a, b)$. By Corollary 3.7 (taking $H' = K_2$ and $\bar{y} = \bar{y}' = (a, b)$), there is a connected graph (G, w, x) with distinguished vertices w and x such that (w, x) is an edge and

$$|\text{Hom}((G, w, x) \rightarrow (H, a, b))| \not\equiv |\text{Hom}((G, w, x) \rightarrow (K_2, a, b))| \pmod{2}. \quad (5)$$

G must be bipartite — otherwise

$$|\text{Hom}((G, w, x) \rightarrow (H, a, b))| = |\text{Hom}((G, w, x) \rightarrow (K_2, a, b))| = 0,$$

contradicting (5). Thus, $|\text{Hom}((G, w, x) \rightarrow (K_2, a, b))| = 1$, so the edge (a, b) of H together with (G, w, x) is an even gadget. \square

LEMMA 5.16. *Suppose that H is a connected, bipartite, square-free graph with more than one edge such that, for some edge (a, b) of H , (H, a, b) is involution-free. Suppose that every vertex of H has odd degree. Then H has a hardness gadget.*

PROOF. By Lemma 5.15, H has an even gadget. Choose an even gadget consisting of an edge (i, s) of H and a connected bipartite graph G with distinguished edge (w, x) so that $N = |\text{Hom}((G, w, x) \rightarrow (H, i, s))|$ is even. Choose the even gadget so that the number of vertices of G is as small as possible. There is a homomorphism from G to the edge (i, s) so $N > 0$. N is even, so G cannot be a single edge.

First, we show that $\deg_G(w) \geq 2$. Suppose, towards a contradiction, that $\deg_G(w) = 1$, i.e., that x is the only neighbour of w in G . If this is the case, then x must have some

neighbour $w' \neq w$, since G is not a single edge. We have

$$\begin{aligned} 0 &\equiv |\text{Hom}((G, w, x) \rightarrow (H, i, s))| \pmod{2} \\ &\equiv |\text{Hom}((G - w, x) \rightarrow (H, s))| \pmod{2} \\ &= \sum_{c \in \Gamma_H(s)} |\text{Hom}((G - w, x, w') \rightarrow (H, s, c))|. \end{aligned}$$

Since every vertex in H has odd degree, the sum has an odd number of terms. Since the total is even, there must be some c such that $|\text{Hom}((G - w, x, w') \rightarrow (H, s, c))|$ is even, contradicting the choice of G . By the same argument, $\deg_G(x) \geq 2$, also.

For any vertex $v \in V(G)$, let

$$C(v) = \{c \in V(H) \mid |\text{Hom}((G, w, x, v) \rightarrow (H, i, s, c))| \text{ is odd}\}.$$

Note that, for any $v \in V(G)$, $|C(v)|$ is even since, otherwise, N would be odd.

We now show that $C(y) \neq \emptyset$ for every $y \in \Gamma_G(x) \setminus \{w\}$. If $C(y) = \emptyset$, then, in particular, $i \notin C(y)$, so $|\text{Hom}((G, w, x, y) \rightarrow (H, i, s, i))|$ is even. But then $|\text{Hom}((G', w, x) \rightarrow (H, i, s))|$ is even, where G' is the graph made from G by identifying the (distinct) vertices w and y and calling the resulting vertex w . This contradicts minimality in the choice of G . Similarly, $C(z) \neq \emptyset$ for every $z \in \Gamma_G(w) \setminus \{x\}$. Choose vertices $y \in \Gamma_G(x) \setminus \{w\}$ and $z \in \Gamma_G(w) \setminus \{x\}$.

Finally, let J be the partially H -labelled graph $(G, \{w \mapsto i, x \mapsto s\})$ and let $G(J_3)$ be the single edge (y, z) and $\tau(J_3) = \emptyset$. We show that $(i, s, (J, y), (J, z), (J_3, y, z))$ is a hardness gadget for H . $\Omega_y = C(y)$ is even and $i \in C(y)$; likewise, $\Omega_z = C(z)$ is even and $s \in C(z)$.

By the choice of J , $\Omega_y \subseteq \Gamma_H(s)$ and $\Omega_z \subseteq \Gamma_H(i)$. For any $o \in \Omega_y - i$ and $x \in \Omega_z - s$, H contains edges (o, s) , (s, i) and (i, x) so it does not contain the edge (o, x) as it is square-free. Therefore, $|\Sigma_{o,s}| = |\Sigma_{i,s}| = |\Sigma_{i,x}| = 1$ and $|\Sigma_{o,x}| = 0$ and we have established all the conditions of Definition 4.1. \square

6. MAIN THEOREM

We have shown that all connected, square-free, involution-free graphs (and some disconnected graphs, too) have hardness gadgets and that $\oplus\text{HOMSTO}H$ is $\oplus\text{P}$ -complete for any involution-free graph that has a hardness gadget. To deal with graphs that have involutions, we use reduction by involutions. As we noted in the introduction, Faben and Jerrum showed that every graph H has a unique (up to isomorphism) involution-free reduction H^* . They also proved [Faben and Jerrum 2015, Theorem 3.4] that for any graph G , $|\text{Hom}(G \rightarrow H)| \equiv |\text{Hom}(G \rightarrow H^*)| \pmod{2}$. Hence, $\oplus\text{HOMSTO}H$ has the same complexity as $\oplus\text{HOMSTO}H^*$.

If H is a tree (as it was for Faben and Jerrum), then its involution-free reduction H^* is connected. However, for general graphs, the fact that H is connected does not imply that H^* is connected.¹ The final result that we need from Faben and Jerrum is [2015, Theorem 6.1], which allows us to deal with disconnected graphs:

LEMMA 6.1. *Let H be an involution-free graph. If H has a component H' for which $\oplus\text{HOMSTO}H'$ is $\oplus\text{P}$ -complete, then $\oplus\text{HOMSTO}H$ is $\oplus\text{P}$ -complete.*

We can now prove our main result.

¹For example, consider non-isomorphic, disjoint, connected, involution-free graphs H_1 and H_2 and let H be a graph made by adding two disjoint paths of the same length from some vertex $x_1 \in H_1$ to some vertex $x_2 \in H_2$. The only involution of this graph exchanges the interior vertices of the two paths, so $H^* = H_1 \cup H_2$, which is disconnected.

THEOREM 1.2. *Let H be a graph whose involution-free reduction H^* is square-free. If H^* has at most one vertex, then $\oplus\text{HOMSTO}H$ is in P; otherwise, $\oplus\text{HOMSTO}H$ is $\oplus\text{P}$ -complete.*

PROOF. As we noted above, $\oplus\text{HOMSTO}H$ has the same complexity as $\oplus\text{HOMSTO}H^*$. If H^* has at most one vertex, then $\oplus\text{HOMSTO}H^*$ is in P: $|\text{Hom}(G \rightarrow H^*)| = 1$ if G has no edges and $\text{Hom}(G \rightarrow H^*) = \emptyset$ if G has an edge. Otherwise, let H^{**} be any component of H^* with more than one vertex. Such a component must exist since, otherwise, H^* would be a graph with at least two vertices and no edges, and any such graph has an involution.

If H^{**} has two or more vertices of even degree, then it has a hardness gadget by Lemma 5.3. If H^{**} has exactly one vertex of even degree, it has a hardness gadget by Lemma 5.7. If the previous cases do not apply, then every vertex of H^{**} must have odd degree. By Lemma 5.1, H^{**} contains a cycle. If it contains an odd cycle, it has a hardness gadget by Lemma 5.13. Otherwise, H^{**} is bipartite. By construction, H^{**} is connected and square-free. Since H^{**} contains a cycle, it has more than one edge. Since it is involution-free, it certainly contains an edge (a, b) so that (H^{**}, a, b) is involution-free. Every vertex of H^{**} has odd degree, so it has a hardness gadget by Lemma 5.16.

We have established that either H^* has at most one vertex, in which case $\oplus\text{HOMSTO}H^*$ and $\oplus\text{HOMSTO}H$ are in P, or that some component H^{**} of H^* has a hardness gadget. In the latter case, $\oplus\text{HOMSTO}H^{**}$ is $\oplus\text{P}$ -complete by Theorem 4.2. $\oplus\text{HOMSTO}H^*$ is $\oplus\text{P}$ -complete by Lemma 6.1, so $\oplus\text{HOMSTO}H$ is $\oplus\text{P}$ -complete. \square

ACKNOWLEDGMENTS

We thank the referees for very useful comments.

REFERENCES

- F. Arends, J. Ouaknine, and C. W. Wampler. 2011. On Searching for Small Kochen–Specker Vector Systems. In *37th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2011), revised papers (LNCS)*, Vol. 6986. Springer-Verlag, 23–34. DOI : http://dx.doi.org/10.1007/978-3-642-25870-1_4
- M. A. Armstrong. 1988. *Groups and Symmetry*. Springer-Verlag.
- A. A. Bulatov and M. Grohe. 2005. The complexity of partition functions. *Theor. Comput. Sci.* 348, 2–3 (2005), 148–186.
- J.-Y. Cai and P. Lu. 2011. Holographic Algorithms: From Art to Science. *J. Comput. Syst. Sci.* 77, 1 (2011), 41–61.
- M. Conforti, G. Cornuéjols, and K. Vušković. 2004. Square-free perfect graphs. *J. Combin. Theory Ser. B* 90, 2 (2004), 257–307. DOI : <http://dx.doi.org/10.1016/j.jctb.2003.08.003>
- M. E. Dyer and C. S. Greenhill. 2000. The Complexity of Counting Graph Homomorphisms. *Random Struct. Algorithms* 17, 3–4 (2000), 260–289.
- J. Faben and M. Jerrum. 2015. The complexity of parity graph homomorphism: an initial investigation. *Theor. Comput.* 11 (2015), 35–57.
- A. Göbel, L. A. Goldberg, and D. Richerby. 2014. The Complexity of Counting Homomorphisms to Cactus Graphs Modulo 2. *ACM T. Comput. Theory* 6, 4 (2014), article 17.
- L. A. Goldberg, M. Grohe, M. Jerrum, and M. Thurley. 2010. A complexity dichotomy for partition functions with mixed signs. *SIAM J. Comput.* 39, 7 (2010), 3336–3402.
- L. A. Goldberg, R. Gysel, and J. Lapinskas. 2014. Approximately Counting Locally-Optimal Structures. *CoRR* abs/1411.6829 (2014).
- L. M. Goldschlager and I. Parberry. 1986. On the construction of parallel computers from various bases of Boolean functions. *Theor. Comput. Sci.* 43 (1986), 43–58.
- H. Guo, S. Huang, P. Lu, and M. Xia. 2011. The Complexity of Weighted Boolean #CSP Modulo k . In *28th International Symposium on Theoretical Aspects of Computer Science (STACS 2011) (Leibniz International Proceedings in Informatics (LIPIcs))*, Vol. 9. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 249–260. DOI : <http://dx.doi.org/10.4230/LIPIcs.STACS.2011.249>
- P. Hell and J. Nešetřil. 1990. On the complexity of H -coloring. *J. Comb. Theory, Ser. B* 48, 1 (1990), 92–110.

- P. Hell and J. Nešetřil. 2004. *Graphs and Homomorphisms*. Oxford University Press.
- L. Lovász. 1967. Operations with Structures. *Acta Math. Acad. Sci. Hungar.* 18, 3–4 (1967), 321–328.
- C. H. Papadimitriou and S. Zachos. 1982. Two remarks on the power of counting. In *Proc. 6th GI-Conference on Theoretical Computer Science*. Springer-Verlag, 269–275.
- T. J. Schaefer. 1978. The Complexity of Satisfiability Problems. In *Proc. 10th Annual ACM Symposium on Theory of Computing (STOC 1978)*. ACM Press, 216–226.
- S. Toda. 1991. PP is as hard as the polynomial-time hierarchy. *SIAM J. Comput.* 20, 5 (1991), 865–877.
- L. G. Valiant. 2006. Accidental Algorithms. In *Proc. 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006)*. IEEE, 509–517.
- M. Xia, P. Zhang, and W. Zhao. 2007. Computational Complexity of Counting Problems on 3-regular planar graphs. *Theoret. Comput. Sci.* 384, 1 (2007), 111–125.

Received —; revised —; accepted —