

Counting Homomorphisms to Cactus Graphs Modulo 2*

Andreas Göbel, Leslie Ann Goldberg, and David Richerby

Department of Computer Science, University of Oxford, Oxford, UK

Abstract

A homomorphism from a graph G to a graph H is a function from $V(G)$ to $V(H)$ that preserves edges. Many combinatorial structures that arise in mathematics and computer science can be represented naturally as graph homomorphisms and as weighted sums of graph homomorphisms. In this paper, we study the complexity of counting homomorphisms modulo 2. The complexity of modular counting was introduced by Papadimitriou and Zachos and it has been pioneered by Valiant who famously introduced a problem for which counting modulo 7 is easy but counting modulo 2 is intractable. Modular counting provides a rich setting in which to study the structure of homomorphism problems. In this case, the structure of the graph H has a big influence on the complexity of the problem. Thus, our approach is graph-theoretic. We give a complete solution for the class of cactus graphs, which are connected graphs in which every edge belongs to at most one cycle. Cactus graphs arise in many applications such as the modelling of wireless sensor networks and the comparison of genomes. We show that, for some cactus graphs H , counting homomorphisms to H modulo 2 can be done in polynomial time. For every other fixed cactus graph H , the problem is complete for the complexity class $\oplus P$ which is a wide complexity class to which every problem in the polynomial hierarchy can be reduced (using randomised reductions). Determining which H lead to tractable problems can be done in polynomial time. Our result builds upon the work of Faben and Jerrum, who gave a dichotomy for the case in which H is a tree.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems


Keywords and phrases modular counting, homomorphisms, cactus graph, graph algorithms

Digital Object Identifier 10.4230/LIPIcs.STACS.2014.350

1 Introduction

A homomorphism from a graph G to a graph H is a function from $V(G)$ to $V(H)$ that preserves edges (i.e., maps every edge of G to some edge of H). Many combinatorial structures arising in mathematics and computer science can be represented naturally as graph homomorphisms. For example, proper q -colourings of a graph G correspond to homomorphisms from G to the q -clique, and independent sets of G correspond to homomorphisms from G to the 2-vertex connected graph with one self-loop (the set of vertices of G mapped to the unlooped vertex is independent). Partition functions in statistical physics such as the Ising, Potts, and hard-core models arise naturally as weighted sums of homomorphisms. See, e.g., [3, 11].

* 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. Some of the initial research was supported by the EPSRC grant EP/I011528/1. A full version with proofs appears in [10].



■ **Figure 1** $\oplus\text{HOMSTO}H_1$ and $\oplus\text{HOMSTO}H_3$ are $\oplus\text{P}$ -complete, but $\oplus\text{HOMSTO}H_2$ is in FP.

We study the complexity of counting homomorphisms modulo 2. For graphs G and H , let $\text{Hom}(G, H)$ be the set of homomorphisms from G to H . For each fixed H , we study the computational problem $\oplus\text{HOMSTO}H$, i.e., computing $|\text{Hom}(G, H)| \bmod 2$, given input G .

The structure of the graph H has a big influence on the complexity of $\oplus\text{HOMSTO}H$. For example, consider the graphs H_1 , H_2 and H_3 depicted in Figure 1. Our result implies that $\oplus\text{HOMSTO}H_1$ is complete for the class $\oplus\text{P}$ (under polynomial-time Turing reductions). H_2 is constructed by moving the top right “bristle” from H_1 down to the bottom right. Under the standard assumption that $\oplus\text{P} \neq \text{FP}$, moving this bristle makes the problem easier – our result implies that $\oplus\text{HOMSTO}H_2$ is solvable in polynomial time. The graph H_3 is constructed by moving the top bristle from left to right in H_2 . This makes the problem hard again – $\oplus\text{HOMSTO}H_3$ is $\oplus\text{P}$ -complete.

The goal of this research is to study the complexity of $\oplus\text{HOMSTO}H$ for every fixed graph H and to determine for which graphs H the problem is in FP, for which it is $\oplus\text{P}$ -complete, and whether there are any H for which the problem has intermediate complexity. In this paper, we give a complete solution to this problem for the class of *cactus graphs*.

A cactus graph is a connected graph in which every edge belongs to at most one cycle. Cactus graphs were first defined by Harary and Uhlenbeck [13] who attributed them to the physicist Husimi and therefore called them *Husimi Trees*. Cactus graphs arise, for example, in the modelling of wireless sensor networks [2] and in the comparison of genomes [18]. Some NP-hard graph problems can be solved in polynomial time on cactus graphs [1].

1.1 The complexity of modular counting

The complexity of modular counting is an interesting topic with some surprising results and we only mention a few highlights here. It is important to note that $\oplus\text{P}$ (first studied in [12, 17]) is a very large complexity class. We treat $\oplus\text{P}$ from the point of view of function computation: it is all problems of the form “compute $f(x) \bmod 2$ ” where computing $f(x)$ is in $\#\text{P}$. $\oplus\text{P}$ is sufficiently powerful that there is randomised polynomial-time reduction [19] from every problem in the polynomial hierarchy to some problem in $\oplus\text{P}$. Thus, under the natural hypothesis that problems in the higher levels of the polynomial hierarchy are not solvable in (randomised) polynomial time, $\oplus\text{P}$ -complete problems are much harder than problems in FP, which is the class of of function-computation problems that are solvable in polynomial time.

The complexity of counting modulo 2 is different from the complexity of decision problems and counting problems. First, consider an NP-complete decision problem. The mod-2

counting version of this problem can be intractable, as you might expect (for example, counting vertex covers or independent sets modulo 2 is \oplus P-complete [20]) but it can also be tractable. As an example, consider counting proper 3-colourings of a graph modulo 2. There are an even number of 3-colourings that use all three colours, since there are six permutations of these colours. There are also an even number of 3-colourings that use exactly two colours, since the colours can be swapped. It is easy to count 1-colourings, so it is easy to count all proper colourings modulo 2. Next, consider a #P-complete counting problem. The mod-2 counting version of this problem can be intractable or tractable, as the examples given above illustrate. As another example where the mod-2 counting version is tractable, consider the problem of computing the permanent of a matrix modulo 2. Since $-1 \equiv 1 \pmod{2}$, the permanent is equal modulo 2 to the determinant, so it can readily be computed in polynomial time.

Another interesting aspect of modular counting is the fact that the value of the modulus can affect the tractability of the problem. As an example, consider the well-known work of Valiant [20] which identified a certain satisfiability problem where satisfying assignments are easy to count modulo 7 but difficult to count modulo 2.

1.2 Dichotomies for graph homomorphism problems

Determining the border between tractability and intractability for large classes of modular counting problems is an important step towards understanding the structure of the problems themselves. In this paper we work within the context of graph homomorphism problems because graph homomorphisms are general enough to capture a wide variety of combinatorial problems, yet they exhibit sufficient structure that dichotomies exist. Hell and Nešetřil [14] pioneered this direction by completely classifying undirected graphs according to the difficulty of the graph homomorphism decision problem. They showed if a fixed graph H has a self-loop, or is bipartite then the problem of determining whether an input graph has a homomorphism to H is in P. For every other fixed graph H , the decision problem is NP-complete.

Over recent years, dichotomy theorems have also been established for the problem of counting graph homomorphisms and computing weighted sums of homomorphisms. Dyer and Greenhill [7] showed that the problem of counting homomorphisms to H is solvable in polynomial time if every component of H is an isolated vertex, a complete graph with all self-loops present, or a complete bipartite graph with no self-loops. For every other H , it is #P-complete. In particular, there are no graphs H for which the problem has intermediate complexity. This dichotomy was extended to the problem of computing weighted sums of homomorphisms to H . A dichotomy was given by Bulatov and Grohe [3] for the case where the weights are positive, by Goldberg, Grohe, Jerrum and Thurley [11] for the case where the weights are real, and by Cai, Chen and Lu [4] for complex weights.

1.3 Counting graph homomorphisms modulo 2

The first results on the complexity of counting graph homomorphisms modulo 2 were obtained by Faben and Jerrum [8, 9], who made some important structural discoveries which we also use.

An *involution* of a graph is an automorphism of order 2. If σ is an automorphism of a graph H then H^σ denotes the subgraph of H induced by the fixed points of σ .

► **Lemma 1.** ([9, Lemma 3.3]) *If H is a graph and σ is an involution of H then, for any graph G , $|\text{Hom}(G, H)| \equiv |\text{Hom}(G, H^\sigma)| \pmod{2}$.*

The lemma is useful because it enables us to reduce the problem of counting homomorphisms to H modulo 2 to the problem of counting homomorphisms to H^σ modulo 2. This leads naturally to the idea of reduction by involutions. Let \rightarrow be the relation on graphs where $H \rightarrow H'$ if and only if there is an involution σ of H such that $H' = H^\sigma$. Let \rightarrow^* be the transitive closure of \rightarrow . Faben and Jerrum showed that repeatedly applying \rightarrow to a graph H reduces H to a unique involution-free graph, up to isomorphism. Also, to classify the complexity of counting homomorphisms to H , it suffices to study the complexity of counting homomorphisms to its connected components.

► **Lemma 2.** ([9, Theorem 6.1]) *Let H be an involution-free graph. If H has a connected component H_1 such that $\oplus\text{HOMSTO}H_1$ is $\oplus\text{P}$ -hard with respect to polynomial-time Turing reductions, then $\oplus\text{HOMSTO}H$ is also $\oplus\text{P}$ -hard.*

It is easy to see that, if $\oplus\text{HOMSTO}H_j$ is solvable in polynomial time for every connected component H_j of H , then $\oplus\text{HOMSTO}H$ is also solvable in polynomial time. Faben and Jerrum used the structural results to give a dichotomy for the complexity of $\oplus\text{HOMSTO}H$ when H is a tree. Define the “involution-free reduction” H' of a graph H to be the lexicographically-minimal involution-free graph such that $H \rightarrow^* H'$. We can state their result as follows.

► **Theorem 3.** ([9, Theorem 3.8]) *If H is a tree then $\oplus\text{HOMSTO}H$ is $\oplus\text{P}$ -complete if the involution-free reduction of H has more than one vertex. Otherwise, it is in FP.*

Every involution-free tree is asymmetric (has no non-trivial automorphisms). Thus, the technical work of proving Theorem 3 is to show that $\oplus\text{HOMSTO}H$ is $\oplus\text{P}$ -hard for every asymmetric tree H with more than one vertex. Fortunately, this can be done without too much technical complexity. Developing a dichotomy to cover all graphs seems to be much harder and even the dichotomy for cactus graphs requires a substantial technical effort, as we will see. Nevertheless, there is a general conjecture as to what the outcome would be.


► **Conjecture 4** (Faben and Jerrum). *Let H be a (not necessarily simple) graph. $\oplus\text{HOMSTO}H$ is in FP if the involution-free reduction of H is empty, a single vertex (with or without a self-loop) or a graph with two isolated vertices, exactly one having a self-loop. Otherwise, it is $\oplus\text{P}$ -complete.*

1.4 Our result

Recall that a cactus graph is a connected, simple graph in which every edge belongs to at most one cycle. Our main result gives a proof of Faben and Jerrum’s conjecture for cactus graphs.

► **Theorem 5.** *Let H be a simple graph with every edge in at most one cycle. If the involution-free reduction of H has at most one vertex, then $\oplus\text{HOMSTO}H$ is solvable in polynomial time. Otherwise, $\oplus\text{HOMSTO}H$ is complete for $\oplus\text{P}$ under polynomial-time Turing reductions.*

To prove this, we must investigate all involution-free cactus graphs, not just the asymmetric ones. This is because, unlike the situation for trees, there are involution-free cactus graphs, such as H_4 in Figure 2, that have non-trivial automorphisms. This graph has no involutions but has an automorphism of order 3 which rotates the cycle. Incidentally, it is easy to see that the graph H_5 in the figure has an involution that moves all vertices, so $\oplus\text{HOMSTO}H_5$ is in FP. Our result implies that $\oplus\text{HOMSTO}H_4$ is $\oplus\text{P}$ -complete.



■ **Figure 2** $\oplus\text{HOMSTO}H_4$ is $\oplus\text{P}$ -complete but $\oplus\text{HOMSTO}H_5$ is in FP.

To prove the hardness result in Theorem 5, we introduce three graph-theoretic notions: hardness gadgets, partial hardness gadgets, and mosaics. Hardness gadgets and partial hardness gadgets are, as the name suggests, structures for proving $\oplus\text{P}$ -hardness. Mosaics are graphs built on unions of 4-cycles. They are what is left in inductive cases where hardness gadgets don't exist and we use them in our inductive proof. Our approach is therefore recursive: we decompose involution-free cactus graphs at cut vertices so that every component contains at least one of these three induced structures. We then combine these structures to obtain hardness gadgets in the original graph. If an asymmetric graph H contains a hardness gadget, then it is relatively easy to show that $\oplus\text{HOMSTO}H$ is $\oplus\text{P}$ -complete — the proof is by reduction from the problem of counting independent sets modulo 2, generalising the argument for trees. We will discuss the situation in which H is not asymmetric presently.

Even when H is asymmetric, the most difficult part of the argument is showing that every non-trivial involution-free cactus graph does actually contain a hardness gadget. The presence of cycles greatly complicates this argument, hence the need to define hardness gadgets, partial hardness gadgets and mosaics and to decompose cactus graphs into components with these three different structures, which can then be combined to form hardness gadgets.

When the graph has non-trivial automorphisms, there is a further complication. Suppose that G and H are graphs and that p is a function from $V(G)$ to $2^{V(H)}$. A homomorphism f from G to H is said to satisfy the “pinning” function p if, for every $v \in V(G)$, we have $f(v) \in p(v)$. Now suppose that H is an involution-free graph containing a hardness-gadget. The high-level strategy for proving that $\oplus\text{HOMSTO}H$ is $\oplus\text{P}$ -hard is to first reduce the problem of counting independent sets modulo 2 to the problem of counting pinned homomorphisms from G to H (modulo 2) and then to reduce the latter problem to $\oplus\text{HOMSTO}H$. This pinning approach has been used successfully in dichotomy theorems in related domains [3, 5, 6]. When H is asymmetric, the application of pinning works smoothly. Building on work of Lovász [15], Faben and Jerrum reduced the pinned problem to the unpinned one for the case in which the pinning function pins some vertex to an orbit in the automorphism group of H . When H is asymmetric (as it is, when H is a tree), the orbit is just a single vertex, and this is just what is required. If H is not asymmetric, we do not know how to pin a vertex of G to a particular vertex in H . To get around this, we augment G with a copy of H and we pin every vertex in the copy to its own orbit in the automorphism group of H . Every homomorphism from an involution-free cactus graph to itself that respects the orbits of all of its vertices is, in fact, an automorphism of H , and this enables us to solve the problem.

Theorem 5 gives a dichotomy for cactus graphs. If the involution-free reduction of H has at most one vertex then $\oplus\text{HOMSTO}H$ is in FP. Otherwise, it is $\oplus\text{P}$ -complete. Furthermore, the meta-problem of determining which is the case, given input H , is computationally easy. Finding an involution of H reduces in polynomial time to computing the size of H 's automorphism group modulo 2. The latter problem is in FP for cactus graphs because, for example, they are planar.

1.5 Notation

Given two graphs G and H (not necessarily vertex-disjoint), $G \cup H$ is the graph $(V(G) \cup V(H), E(G) \cup E(H))$. If E is a set of edges, let $V(E)$ denote the set of endpoints of edges in E and let $G \cup E$ denote the graph $G \cup (V(E), E)$. Given a set $V' \subseteq V(G)$, let $G - V' = G[V(G) \setminus V']$. We use the phrase “ j -walk” in a graph to refer to a walk of length j .

We use $\Gamma_H(v)$ to denote the set of neighbours of vertex v in H . A *rooted graph* is a pair (H, x) where H is a graph and $x \in V(H)$ is a distinguished vertex, the *root*. An automorphism of (H, x) is an automorphism of H that fixes x .

We use $\text{Aut}(H)$ to denote the automorphism group of H and, for $v \in V(H)$, we use $\text{Orb}_H(v)$ to denote the set of vertices of H in the orbit of v under the action of $\text{Aut}(H)$.

2 Pinning, gadgets and mosaics

In this section, we discuss pinning and define the gadgets we use to prove $\oplus\text{P}$ -hardness of $\oplus\text{HOMSTO}H$ problems by reduction from $\oplus\text{IS}$, counting independent sets modulo 2.

Recall from the introduction that a homomorphism $f: V(G) \rightarrow V(H)$ satisfies a *pinning function* $p: V(G) \rightarrow 2^{V(H)}$ if $f(v) \in p(v)$ for all $v \in V(G)$. Let $\text{HomPin}(G, H, p)$ be the set of homomorphisms from G to H that satisfy the pinning function p . Say that a pinning p is r -restrictive if at most r vertices $v \in V(G)$ have $p(v) \neq V(H)$ and for each such vertex v , $p(v)$ is a union of orbits of the automorphism group of H . We consider the following computational problem, which is parameterised by a graph H and a natural number r .

Name: $\oplus r$ -PINNEDHOMSTO H .

Input: A graph G and a r -restrictive pinning function $p: V(G) \rightarrow 2^{V(H)}$.

Output: $|\text{HomPin}(G, H, p)| \pmod{2}$.

Extending the work of Faben and Jerrum who, in turn, built on results of Lovász [15], we prove the following theorem.

► **Theorem 6.** *Let H be an involution-free graph and let r be a positive integer. There is a polynomial-time Turing reduction from $\oplus r$ -PINNEDHOMSTO H to $\oplus\text{HOMSTO}H$.*

We next introduce machinery that we will use to prove that $\oplus r$ -PINNEDHOMSTO H is $\oplus\text{P}$ -complete when H is an involution-free cactus graph and r is defined appropriately.

► **Definition 7.** A *hardness gadget* in a graph H is a tuple $(\beta, s, t, O, i, K, k, w)$ where β is a positive integer, s, t and i are vertices of H , $(O, \{i\}, K)$ is a partition of $\Gamma_H(s)$, and $k: K \rightarrow \mathbb{N}_{>0}$ and $w: K \rightarrow V(H)$ are functions. The following conditions must be satisfied.

1. $|O|$ is odd.
2. For any $o \in O$ and $y \in O \cup \{i\}$, s is the unique vertex that is adjacent to o and y and has an odd number of β -walks to t .
3. There are an even number of $(1 + \beta)$ -walks from i to t .
4. For all $u \in K$, $w(u)$ has an even number of $k(u)$ -walks to u and an odd number of $k(u)$ -walks to every vertex in $O \cup \{i\}$.

These conditions simplify if $\beta = 1$, since having an odd number of 1-walks to a vertex is the same as being adjacent to it.

The construction used in our reduction from $\oplus\text{IS}$ is given formally in Definition 12. Given a graph G and a hardness gadget Γ , we will produce a graph G_Γ that includes a copy of $V(G)$. We call the vertices in this copy, “ G -vertices”. We will use pinning to consider homomorphisms from G_Γ to H that map all G -vertices to neighbours of s . Part 4 of Definition 7 ensures that there will be an even number of such homomorphisms that map any

G -vertices to members of K . These contribute nothing to the total modulo 2 so the effect is to restrict to homomorphisms that map every G -vertex to $O \cup \{i\}$. Part 3 of the definition will ensure that the number of homomorphisms that map adjacent vertices in G to i is even, so these also do not contribute. Thus, the homomorphisms that remain are those in which an independent set of G -vertices are mapped to i . Our key technical result is that every non-trivial, involution-free cactus graph contains a hardness gadget (Theorem 10).

In some cases, our decomposition might yield subgraphs that do not contain hardness gadgets. We are still able to make progress using structures that can be combined with other parts of the graph to produce a hardness gadget. A partial hardness gadget is, essentially, a simplified hardness gadget that has $K = \emptyset$ and that doesn't yet have a "t" vertex: at a later point, we will find a vertex t with the properties necessary to produce a full hardness gadget.

► **Definition 8.** A *partial hardness gadget* in a rooted graph (H, x) is a tuple (s, i, O, P) , where s is a vertex of H , $(\{i\}, O)$ is a partition of $\Gamma_H(s)$, and P is a path in H . The tuple satisfies the following conditions.

1. $|O|$ is odd.
2. P is the unique shortest path from x to i in H .
3. Ps is the unique shortest path from x to s in H .
4. For each $o \in O$, Pso is the unique shortest path from x to o in H .

The final structures arising in our decompositions are "mosaics". Some of these (those with "shortcuts", defined below) already contain hardness gadgets. In other cases, a mosaic will provide a "t" vertex for a partial hardness gadget elsewhere in the decomposed graph.

► **Definition 9.** An unbristled mosaic is the one-vertex rooted graph or a rooted cactus graph that is a union of 4-cycles. A *mosaic* is a rooted graph (H, x) for which there is a partition (V', V'') of $V(H)$ such that: $x \in V'$, $(H[V'], x)$ is an unbristled mosaic, and $E(H) \setminus E(H[V'])$ is a matching between V'' and a subset of V' . The edges of the matching are called *bristles*.

The graphs in Figure 1 would be mosaics if a root were placed at any vertex on a cycle. Note that every vertex of a mosaic is adjacent to at most one bristle, and that the one-vertex rooted graph and a rooted edge are both mosaics.


A *shortcut* in a mosaic (H, x) is a pair of odd-degree vertices, with degree at least 3, that have a unique shortest path P between them, and this path does not contain x . In the full paper, we show that every mosaic with a shortcut contains a hardness gadget.

3 Finding hardness gadgets

In Sections 6 and 7 of the full paper, we prove the following result.

► **Theorem 10.** *Every involution-free cactus graph H with more than one vertex contains a hardness gadget.*

Given a cut vertex v of a graph H , let H'_1, \dots, H'_κ be the connected components of $H - \{v\}$. Let the *split* of H at v be the set of graphs $\{H_1, \dots, H_\kappa\}$, where $H_j = H[V(H'_j) \cup \{v\}]$. To prove Theorem 10, we mostly proceed by splitting at cut vertices and investigating the resulting components. A key point is that, if $\{H_1, \dots, H_\kappa\}$ is the split of an involution-free graph H at a cut vertex v then each rooted graph (H_j, v) is involution-free, even though the unrooted graph H_j might not be. This allows us to perform an induction on rooted graphs to establish the following lemma. Theorem 10 then follows by choosing an appropriate root and constructing a hardness gadget from the contents of the split at the root.



■ **Figure 3** An example graph illustrating the proof ideas of Theorem 10 and Lemma 11.

► **Lemma 11.** *Every involution-free rooted cactus graph (H, x) contains a hardness gadget, contains a partial hardness gadget or is a shortcut-free mosaic.*

Rather than attempting to sketch the lengthy and technical proof of Lemma 11, we will work through an example that illustrates the main techniques. Consider the cactus graph of Figure 3. It is involution-free (in fact, asymmetric) and its split at the vertex x gives the three involution-free rooted graphs (H_6, x) , (H_7, x) and (H_8, x) .

We see immediately that (H_6, x) is a mosaic, and it is shortcut-free, since it has only one odd-degree vertex on a cycle (the degree of x in H_6 is two). Note also that (H_6, x) is asymmetric but the unrooted graph H_6 has an involution that exchanges x with the vertex at distance 2 from it on the same 4-cycle.

Consider, now, (H_7, x) . This graph contains the partial hardness gadget $(s_7, i_7, \{o_7\}, xzi_7)$: $(\{i_7\}, \{o_7\})$ partitions $\Gamma_{H_7}(s_7)$, $|\{o_7\}|$ is odd, xzi_7 is the unique shortest $x-i_7$ path, xzi_7s_7 is the unique shortest $x-s_7$ path and $xzi_7s_7o_7$ is the unique shortest $x-o_7$ path.

Now, we turn our attention to (H_8, x) , in which we will demonstrate a hardness gadget. In the first instance, consider the graph without the dashed path, the easier case. As the notation suggests, we take $s = s_8$ and $i = i_8$. A helpful feature for us here is the even-length cycle that includes these two vertices: by choosing t to be the vertex t_8 , half way around the cycle from i , and taking $\beta = 2$ (so the length of the cycle is $2(\beta + 1)$), we ensure that requirement 3 of the definition of hardness gadgets is met (an even number of $(\beta + 1)$ -walks from i to t). We take $O = \{o_8, o'_8, o''_8\}$, which has odd cardinality so satisfies requirement 1. Requirement 2 is that, for each $o \in O$ and $y \in O \cup \{i\}$, s is the unique vertex adjacent to o and y that has an odd number of β -walks to t . $\beta = 2$ and s and x are the only vertices that send an odd number of 2-walks to t . Since x is not adjacent to any vertex in O , s meets the requirement. Finally, since $(O, \{i\})$ is already a partition of $\Gamma_{H_8}(s)$, we set $K = \emptyset$ and requirement 4 is vacuous. Therefore, writing \perp for the function with empty domain,

$$\Gamma = (\beta, s, t, O, i, K, k, w) = (2, s_8, t_8, \{o_8, o'_8, o''_8\}, i_8, \emptyset, \perp, \perp)$$

is a hardness gadget in H_8 .

To demonstrate a hardness gadget with non-empty K , consider the rooted cactus graph (H'_8, x) formed by adding the dashed edges to H_8 . We take s, t, O, i and β as before but, now, $(O, \{i\})$ is not a partition of $\Gamma_{H'_8}(s)$. Thus, we set $K = \{u\}$, $k(u) = 2$ and $w(u) = u$. u has two 2-walks to itself and one to i and each vertex in O , so requirement 4 is met.

Let us recap: we have split the graph H at cut vertex x and demonstrated that each of the three components of this split contains a hardness gadget or a partial hardness gadget,

or is a shortcut-free mosaic. We now illustrate Theorem 10 by showing how to combine these to produce a hardness gadget in H .

In fact, this is rather easy because the hardness gadget Γ in H_8 is also a hardness gadget in H . This is because the requirements for being a hardness gadget depend on the number of 3-walks from i_8 , o_8 , o'_8 and o''_8 to t_8 and the number of 2-walks from u to vertices adjacent to s_8 ; however, none of these walks can ever leave H_8 . In the full version of the paper, we give formal *distance requirements* that allow us to determine more generally when a hardness gadget in an induced subgraph of H is also a hardness gadget in H .

Our goal is to illustrate the proof techniques, so we will continue and find a second hardness gadget in H by combining the mosaic (H_6, x) with the partial hardness gadget $(s_7, i_7, \{o_7\}, xzi_7)$ in (H_7, x) . As we remarked earlier, if we can find appropriate values for t and β , the partial hardness gadget will become a hardness gadget with $K = \emptyset$. The properties we require of t and β are the following:

- there are an even number of $(1 + \beta)$ -walks from i_7 to t ;
- s_7 is the unique vertex adjacent to o_7 that has an odd number of β -walks to t ; and
- s_7 is the unique vertex adjacent to both o_7 and i_7 that has an odd number of β -walks to t .

Since s_7 is the only vertex adjacent to both o_7 and i_7 , the third property follows from the second. To make the second property easy to verify, we will choose t to have a unique shortest path in H to o_7 , and this path will go through s and have length $1 + \beta$.

Consider the vertices t_6 and t'_6 , which are not adjacent but are on the same cycle, and which have degree 2 and 3, respectively. Further, each has a unique shortest path to x and these two paths differ only in their last edge. It is not hard to see that every involution-free mosaic with at least one cycle must contain a pair of vertices with these properties and, in the full paper, we call such a pair of vertices, along with the shared section of their shortest paths to x , a *2,3-path*.


We are going to take $\beta = 6$ (the distance from s_7 to $\{t_6, t'_6\}$) and we claim that we can choose one of $t = t_6$ or $t = t'_6$ to satisfy the first two properties. In fact, either choice satisfies the second property, since either choice for t gives a unique 7-walk to o_7 .

To verify the claim, we will show that t_6 and t'_6 have different numbers of 7-walks to i_7 , modulo 2. Therefore, one of them has an even number of 7-walks, and that will be our choice for t . There is a unique 5-path from i_7 to each of t_6 and t'_6 : write this path as $x_1x_2 \dots x_6$, where $i_7 = x_1$. Every 7-walk from x_1 to x_6 is of one of the following two types:

1. walks that replace one of the edges (x_3, x_4) , (x_4, x_5) or (x_5, x_6) by going along the other three edges of the 4-cycle that contains it; and
2. walks that replace one of the vertices x_a ($1 \leq a \leq 6$) with the 2-walk x_ayx_a , for some $y \in \Gamma_H(x_a)$.

There are exactly three type-1 walks from i_7 to each of t_6 and t'_6 . The number of type-2 walks from i_7 to t'_6 is exactly one greater than the number to t_6 . For $1 \leq a \leq 5$, there are the same number of choices for y in each case; however, for $a = 6$, there are three choices of y from t'_6 but only two from t_6 . Therefore, the number of 7-walks from i_7 to exactly one of t_6 and t'_6 is even, and we choose that vertex to be t . The reader is invited to check that there are, in total, twenty 7-walks to t'_6 and nineteen to t_6 . Thus, the hardness gadget is

$$(\beta, s, t, O, i, K, k, w) = (6, s_7, t'_6, \{o_7\}, i_7, \emptyset, \perp, \perp).$$



■ **Figure 4** The induced subgraph of G_Γ corresponding to the edge $(x, y) \in E(G)$, with $K = \{u_1, \dots, u_j\}$. H -vertices have double circles and are pinned in the proof of Theorem 13.

4 Counting homomorphisms to cactus graphs

Having shown that every involution-free cactus graph with more than one vertex contains a hardness gadget, we now use these gadgets to show $\oplus P$ -completeness of $\oplus \text{HOMSTO}H$ for non-trivial involution-free cactus graphs H . The reduction is from $\oplus \text{IS}$, which is $\oplus P$ -complete [20]. The reduction is more complicated than the case for trees because an involution-free cactus graph is not necessarily asymmetric — recall the graph H_4 in Figure 2.

In the following definition, “adding a new path P from x to y ” in a graph G means forming a graph $G \cup P$ where $V(G) \cap V(P) = \{x, y\}$.

► **Definition 12.** Let $\Gamma = (\beta, s, t, O, i, K, k, w)$ be a hardness gadget in H . For any graph G , we construct the graph G_Γ as follows. Begin with the graph $G' = (V', E(H))$ where $V' = V(G) \cup V(H) \cup \{v_e \mid e \in E(G)\}$ (these three sets are assumed to be disjoint) and add:

- for every vertex $x \in V(G)$, the edge (x, s) ;
- for every edge $e = (x, y) \in E(G)$, the edges (x, v_e) and (y, v_e) ;
- for every edge $e \in E(G)$, a new β -path $P_{t,e}$ from t to v_e ; and
- for every vertex $x \in V(G)$ and every $u \in K$, a new $k(u)$ -path $P_{x,u}$ from x to $w(u)$.

In G_Γ , we refer to vertices that are in $V(G)$ as G -vertices and those in $V(H)$ as H -vertices. Figure 4 illustrates the construction.

Our construction, G_Γ , is more complex than the construction used for trees, because our hardness gadgets are more general than the corresponding structures in trees and because we must deal with graphs H that are involution-free but still have non-trivial automorphisms. To see the problem of non-trivial automorphisms, consider an involution-free cactus graph H that contains a hardness gadget Γ that is moved by an automorphism π of H . We want to pin one vertex to the s -vertex of Γ and another to the t -vertex. However, we can only pin to the orbits of these vertices, which include $\pi(s)$ and $\pi(t)$, respectively. We must avoid counting “inconsistent” homomorphisms that, for example, map the first vertex to s and the second to $\pi(t)$ because we do not know how many of these homomorphisms exist.

► **Theorem 13.** $\oplus \text{HOMSTO}H$ is $\oplus P$ -complete for every involution-free cactus graph H that contains a hardness gadget.

Proof (sketch). Using Theorem 6, it suffices to reduce $\oplus \text{IS}$ to $\oplus r$ -PINNEDHOMSTO H where $r = |V(H)|$. Let G be the graph whose independent sets we wish to count and let $\Gamma = (\beta, s, t, O, i, K, k, w)$ be a hardness gadget in H . Let p be the pinning function that maps every H -vertex v to $\text{Orb}_H(v)$ and every other vertex of G_Γ to $V(H)$ and let Φ be the set of

homomorphisms from G_Γ to H that satisfy p . Let $\mathcal{I}(G)$ be the set of independent sets in G . We claim that $|\Phi| \equiv |\mathcal{I}(G)| \pmod{2}$.

It can be shown that any $\phi \in \Phi$ acts as an automorphism on the H -vertices. Let $\Phi_\pi \subseteq \Phi$ be set of homomorphisms where this automorphism is π . Writing id for the trivial automorphism, every $\phi \in \Phi_{\text{id}}$ has $\phi(s) = s$ and, for all G -vertices v , $\phi(v) \in O \cup \{i\} \cup K$. For each G -vertex v , $|\{\phi \in \Phi_{\text{id}} \mid \phi(v) \in K\}|$ is even because, when $\phi(v) \in K$, $P_{v,w(\phi(v))}$ can map to an even number of $k(\phi(v))$ -walks in H . So, to compute $|\Phi_{\text{id}}|$ modulo 2, it suffices to count the homomorphisms $\phi \in \Phi_{\text{id}}$ where $\phi(v) \in O \cup \{i\}$ for all G -vertices v . For such a homomorphism, let S_ϕ be the set of G -vertices mapped to i . If $S \in \mathcal{I}(G)$, each G -vertex not in S can map to any of the odd number of elements of O and, for each edge e , we must have $\phi(v_e) = s$ and $P_{t,e}$ can map to an odd number of β -walks in H . If $S \notin \mathcal{I}(G)$, there are an even number of homomorphisms ϕ with $S_\phi = S$ because, if adjacent G -vertices x and y map to i , there are an even number of ways to map the paths $P_{t,(x,y)}x$ and $P_{t,(x,y)}y$ to $(1 + \beta)$ -walks in H . Thus, $|\Phi_{\text{id}}| \equiv |\mathcal{I}(G)| \pmod{2}$. For any automorphism π of H , $|\Phi_\pi| = |\Phi_{\text{id}}|$, so $|\Phi| = |\Phi_{\text{id}}| |\text{Aut } H|$. H is involution-free so, by Cauchy's Group Theorem [16], $|\text{Aut}(H)|$ is odd, so $|\Phi| \equiv |\Phi_{\text{id}}| \equiv |\mathcal{I}(G)| \pmod{2}$. ◀

We can now prove our main result.

► **Theorem 5.** *Let H be a simple graph with every edge in at most one cycle. If the involution-free reduction of H has at most one vertex, then $\oplus\text{HOMSTO}H$ is solvable in polynomial time. Otherwise, $\oplus\text{HOMSTO}H$ is complete for $\oplus\text{P}$ under polynomial-time Turing reductions.*

Proof. Let H' be the involution-free reduction of H . If H' has at most one vertex then $\oplus\text{HOMSTO}H'$ is trivially solvable in polynomial time. By Lemma 1, every graph G satisfies $|\text{Hom}(G, H)| \equiv |\text{Hom}(G, H')| \pmod{2}$ so $\oplus\text{HOMSTO}H$ is also solvable in polynomial time.

If H' has more than one vertex, then some component H_1 of H' has more than one vertex (since H' is involution-free). Also, H_1 is involution-free. Since H_1 is an induced subgraph of H , it is a cactus graph. By Theorems 10 and 13, $\oplus\text{HOMSTO}H_1$ is $\oplus\text{P}$ -hard. By Lemma 2, $\oplus\text{HOMSTO}H'$ is $\oplus\text{P}$ -hard. But Lemma 1 gives a reduction from $\oplus\text{HOMSTO}H'$ to $\oplus\text{HOMSTO}H$, so $\oplus\text{HOMSTO}H$ is also $\oplus\text{P}$ -hard. ◀

References

- 1 B. Ben-Moshe, B. K. Bhattacharya, Q. Shi, and A. Tamir. Efficient algorithms for center problems in cactus networks. *Theor. Comput. Sci.*, 378(3):237–252, 2007.
- 2 B. Ben-Moshe, A. Dvir, M. Segal, and A. Tamir. Centdian computation in cactus graphs. *J. Graph Algorithms Appl.*, 16(2):199–224, 2012.
- 3 A. A. Bulatov and M. Grohe. The complexity of partition functions. *Theor. Comput. Sci.*, 348(2–3):148–186, 2005.
- 4 J.-Y. Cai, X. Chen, and P. Lu. Graph homomorphisms with complex values: A dichotomy theorem. In *Proc. ICALP (1)*, pages 275–286, 2010.
- 5 N. Creignou and M. Hermann. Complexity of generalized satisfiability counting problems. *Inform. Comput.*, 125(1):1–12, 1996.
- 6 M. E. Dyer, L. A. Goldberg, and M. Jerrum. The complexity of weighted Boolean $\#\text{CSP}$. *SIAM J. Comput.*, 38(5):1970–1986, 2009.
- 7 M. E. Dyer and C. S. Greenhill. The complexity of counting graph homomorphisms. *Random Struct. Algorithms*, 17(3–4):260–289, 2000.
- 8 J. Faben. *The Complexity of Modular Counting in Constraint Satisfaction Problems*. PhD thesis, Queen Mary, University of London, 2012.

- 9 J. Faben and M. Jerrum. The complexity of parity graph homomorphism: an initial investigation. *CoRR*, abs/1309.4033, 2013.
- 10 Andreas Göbel, Leslie Ann Goldberg, and David Richerby. The complexity of counting homomorphisms to cactus graphs modulo 2. *CoRR*, abs/1307.0556, 2013.
- 11 L. A. Goldberg, M. Grohe, M. Jerrum, and M. Thurley. A complexity dichotomy for partition functions with mixed signs. *SIAM J. Comput.*, 39(7):3336–3402, 2010.
- 12 L. M. Goldschlager and I. Parberry. On the construction of parallel computers from various bases of Boolean functions. *Theor. Comput. Sci.*, 43:43–58, 1986.
- 13 F. Harary and G. E. Uhlenbeck. On the number of Husimi trees. I. *Proc. Nat. Acad. Sci. U. S. A.*, 39:315–322, 1953.
- 14 P. Hell and J. Nešetřil. On the complexity of H -coloring. *J. Comb. Theory, Ser. B*, 48(1):92–110, 1990.
- 15 L. Lovász. Operations with structures. *Acta Math. Acad. Sci. Hungar.*, 18:321–328, 1967.
- 16 J. H. McKay. Another proof of Cauchy’s group theorem. *The American Mathematical Monthly*, 66:119, 1959.
- 17 Christos H. Papadimitriou and Stathis Zachos. Two remarks on the power of counting. In *Proceedings of the 6th GI-Conference on Theoretical Computer Science*, pages 269–276, London, UK, UK, 1982. Springer-Verlag.
- 18 B. Paten, M. Diekhans, D. Earl, J. St. John, J. Ma, B. B. Suh, and D. Haussler. Cactus graphs for genome comparisons. *J. Comput. Biol.*, 18(3):469–481, 2011.
- 19 S. Toda. PP is as hard as the polynomial-time hierarchy. *SIAM J. Comput.*, 20(5):865–877, 1991.
- 20 L. G. Valiant. Accidental algorithms. In *Proc. FOCS*, pages 509–517, 2006.