

HOLANT PROBLEMS FOR REGULAR GRAPHS WITH COMPLEX EDGE FUNCTIONS

MICHAEL KOWALCZYK¹ AND JIN-YI CAI²

¹ Department of Mathematics and Computer Science, Northern Michigan University
Marquette, MI 49855, USA
E-mail address: mkowalcz@nmu.edu

² Computer Sciences Department, University of Wisconsin, Madison, WI 53706, USA
E-mail address: jyc@cs.wisc.edu

ABSTRACT. We prove a complexity dichotomy theorem for Holant Problems on 3-regular graphs with an arbitrary complex-valued edge function. Three new techniques are introduced: (1) higher dimensional iterations in interpolation; (2) Eigenvalue Shifted Pairs, which allow us to prove that a pair of combinatorial gadgets *in combination* succeed in proving #P-hardness; and (3) algebraic symmetrization, which significantly lowers the *symbolic complexity* of the proof for computational complexity. With *holographic reductions* the classification theorem also applies to problems beyond the basic model.

1. Introduction

In this paper we consider the following subclass of Holant Problems [5, 6]. An input regular graph $G = (V, E)$ is given, where every $e \in E$ is labeled with a (symmetric) edge function g . The function g takes 0-1 inputs from its incident nodes and outputs arbitrary values in \mathbb{C} . The problem is to compute the quantity $\text{Holant}(G) = \sum_{\sigma: V \rightarrow \{0,1\}} \prod_{\{u,v\} \in E} g(\{\sigma(u), \sigma(v)\})$.

Holant Problems are a natural class of counting problems. As introduced in [5, 6], the general Holant Problem framework can encode all Counting Constraint Satisfaction Problems (#CSP). This includes special cases such as weighted VERTEX COVER, GRAPH COLORINGS, MATCHINGS, and PERFECT MATCHINGS. The subclass of Holant Problems in this paper can also be considered as (weighted) H -homomorphism (or H -coloring) problems [2, 3, 7, 8, 9, 10] with an arbitrary 2×2 symmetric complex matrix H , however *restricted* to regular graphs G as input. E.g., VERTEX COVER is the case when $H = \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix}$. When the matrix H is a 0-1 matrix, it is called unweighted. Dichotomy theorems (i.e., the problem is either in P or #P-hard, depending on H) for unweighted H -homomorphisms with undirected graphs H and directed acyclic graphs H are given in [8] and [7] respectively. A

1998 ACM Subject Classification: F.2.1.

Key words and phrases: Computational complexity.

The second author is supported by NSF CCF-0830488 and CCF-0914969.



dichotomy theorem for any symmetric matrix H with non-negative real entries is proved in [2]. Goldberg et al. [9] proved a dichotomy theorem for all real symmetric matrices H . Finally, Cai, Chen, and Lu have proved a dichotomy theorem for all complex symmetric matrices H [3].

The crucial difference between Holant Problems and #CSP is that in #CSP, EQUALITY functions of arbitrary arity are *presumed* to be present. In terms of H -homomorphism problems, this means that the input graph is allowed to have vertices of arbitrarily high degrees. This may appear to be a minor distinction; in fact it has a major impact on complexity. It turns out that if EQUALITY gates of arbitrary arity are freely available in possible inputs then it is technically easier to prove #P-hardness. Proofs of previous dichotomy theorems make extensive use of constructions called thickening and stretching. These constructions require the availability of EQUALITY gates of arbitrary arity (equivalently, vertices of arbitrarily high degrees) to carry out. Proving #P-hardness becomes more challenging in the degree restricted case. Furthermore there are indeed cases within this class of counting problems where the problem is #P-hard for general graphs, but solvable in P when restricted to 3-regular graphs.

We denote the (symmetric) edge function g by $[x, y, z]$, where $x = g(0, 0)$, $y = g(0, 1) = g(1, 0)$ and $z = g(1, 1)$. Functions will also be called gates or signatures. (For VERTEX COVER, the function corresponding to H is the OR gate, and is denoted by the signature $[0, 1, 1]$.) In this paper we give a dichotomy theorem for the complexity of Holant Problems on 3-regular graphs with arbitrary signature $g = [x, y, z]$, where $x, y, z \in \mathbb{C}$. First, if $y = 0$, the Holant Problem is easily solvable in P. Assuming $y \neq 0$ we may normalize g and assume $y = 1$. Our main theorem is as follows:

Theorem 1.1. *Suppose $a, b \in \mathbb{C}$, and let $X = ab$, $Z = (\frac{a^3+b^3}{2})^2$. Then the Holant Problem on 3-regular graphs with $g = [a, 1, b]$ is #P-hard except in the following cases, for which the problem is in P.*

- (1) $X = 1$.
- (2) $X = Z = 0$.
- (3) $X = -1$ and $Z = 0$.
- (4) $X = -1$ and $Z = -1$.

*If we restrict the input to planar 3-regular graphs, then these four categories are solvable in P, as well as a fifth category $X^3 = Z$. The problem remains #P-hard in all other cases.*¹ ■

These results can be extended to k -regular graphs (we detail how this is accomplished in a forthcoming work). One can also use holographic reductions [15] to extend this theorem to more general Holant Problems.

In order to achieve this result, some new proof techniques are introduced. To discuss this we first take a look at some previous results. Valiant [13, 14] introduced the powerful technique of *interpolation*, which was further developed by many others. In [5] a dichotomy theorem is proved for the case when g is a Boolean function. The technique from [5] is to provide certain algebraic criteria which ensure that *interpolation* succeeds, and then apply these criteria to prove that (a large number yet) finitely many individual problems are #P-hard. This involves (a small number of) gadget constructions, and the algebraic criteria

¹Technically, computational complexity involving complex or real numbers should, in the Turing model, be restricted to computable numbers. In other models such as the Blum-Shub-Smale model [1] no such restrictions are needed. Our results are not sensitive to the exact model of computation.

are powerful enough to show that they succeed in each case. Nonetheless this involves a case-by-case verification. In [6] this theorem is extended to all real-valued a and b , and we have to deal with infinitely many problems. So instead of focusing on one problem, we devised (a large number of) recursive gadgets and analyzed the regions of $(a, b) \in \mathbb{R}^2$ where they fail to prove #P-hardness. The algebraic criteria from [5] are not suitable (Galois theoretic) for general a and b , and so we formulated weaker but simpler criteria. Using these criteria, the analysis of the failure set becomes expressible as containment of semi-algebraic sets. As semi-algebraic sets are decidable, this offers the ultimate possibility that *if* we found enough gadgets to prove #P-hardness, *then* there is a *computational* proof (of computational intractability) in a finite number of steps. However this turned out to be a tremendous undertaking in symbolic computation, and many additional ideas were needed to finally carry out this plan. In particular, it would seem hopeless to extend that approach to all complex a and b .

In this paper, we introduce three new ideas. (1) We introduce a method to construct gadgets that carry out iterations at a higher dimension, and then collapse to a lower dimension for the purpose of constructing unary signatures. This involves a starter gadget, a recursive iteration gadget, and a finisher gadget. We prove a lemma that guarantees that among polynomially many iterations, some subset of them satisfies properties sufficient for interpolation to succeed (it may not be known *a priori* which subset worked, but that does not matter). (2) Eigenvalue Shifted Pairs are coupled pairs of gadgets whose transition matrices differ by δI where $\delta \neq 0$. They have shifted eigenvalues, and by analyzing their failure conditions, we can show that except on very rare points, one or the other gadget succeeds. (3) Algebraic symmetrization. We derive a new expression of the Holant polynomial over 3-regular graphs, with a crucially reduced degree. This simplification of the Holant and related polynomials condenses the problem of proving #P-hardness to the point where all remaining cases can be handled by symbolic computation. We also use the same expression to prove tractability.

The rest of this paper is organized as follows. In Section 2 we discuss notation and background information. In Section 3 we cover interpolation techniques, including how to collapse higher dimensional iterations to interpolate unary signatures. In Section 4 we show how to perform algebraic symmetrization of the Holant, and introduce Eigenvalue Shifted Pairs (ESP) of gadgets. Then we combine the new techniques to prove Theorem 1.1. Some proofs are omitted due to space limitations; a full version will appear in [11].

2. Notations and Background

We state the counting framework more formally. A *signature grid* $\Omega = (G, \mathcal{F}, \pi)$ consists of a labeled graph $G = (V, E)$ where π labels each vertex $v \in V$ with a function $f_v \in \mathcal{F}$. We consider all edge assignments $\xi : E \rightarrow \{0, 1\}$; f_v takes inputs from its incident edges $E(v)$ at v and outputs values in \mathbb{C} . The counting problem on the instance Ω is to compute²

$$\text{Holant}_\Omega = \sum_{\xi} \prod_{v \in V} f_v(\xi|_{E(v)}).$$

Suppose G is a bipartite graph (U, V, E) such that each $u \in U$ has degree 2. Furthermore suppose each $v \in V$ is labeled by an EQUALITY gate $=_k$ where $k = \deg(v)$. Then any non-zero term in Holant_Ω corresponds to a 0-1 assignment $\sigma : V \rightarrow \{0, 1\}$. In fact, we can merge

²The term Holant was first introduced by Valiant in [15] to denote a related exponential sum.

the two incident edges at $u \in U$ into one edge e_u , and label this edge e_u by the function f_u . This gives an edge-labeled graph (V, E') where $E' = \{e_u : u \in U\}$. For an edge-labeled graph (V, E') where $e \in E'$ has label g_e , $\text{Holant}_\Omega = \sum_{\sigma: V \rightarrow \{0,1\}} \prod_{e=(v,w) \in E'} g_e(\sigma(v), \sigma(w))$. If each g_e is the same function g (but assignments $\sigma : V \rightarrow [q]$ take values in a finite set $[q]$) this is exactly the H -coloring problem (for undirected graphs g is a symmetric function). In particular, if (U, V, E) is a $(2, k)$ -regular bipartite graph, equivalently $G' = (V, E')$ is a k -regular graph, then this is the H -coloring problem restricted to k -regular graphs. In this paper we will discuss 3-regular graphs, where each g_e is the same symmetric complex-valued function. We also remark that for general bipartite graphs (U, V, E) , giving EQUALITY (of various arities) to all vertices on one side V defines #CSP as a special case of Holant Problems. But whether EQUALITY of various arities are present has a major impact on complexity, thus Holant Problems are a refinement of #CSP.

A symmetric function $g : \{0, 1\}^k \rightarrow \mathbb{C}$ can be denoted as $[g_0, g_1, \dots, g_k]$, where g_i is the value of g on inputs of Hamming weight i . They are also called *signatures*. Frequently we will revert back to the bipartite view: for $(2, 3)$ -regular bipartite graphs (U, V, E) , if every $u \in U$ is labeled $g = [g_0, g_1, g_2]$ and every $v \in V$ is labeled $r = [r_0, r_1, r_2, r_3]$, then we also use $\#[g_0, g_1, g_2] \mid [r_0, r_1, r_2, r_3]$ to denote the Holant Problem. Note that $[1, 0, 1]$ and $[1, 0, 0, 1]$ are EQUALITY gates $=_2$ and $=_3$ respectively, and the main dichotomy theorem in this paper is about $\#[x, y, z] \mid [1, 0, 0, 1]$, for all $x, y, z \in \mathbb{C}$. We will also denote $\text{Hol}(a, b) = \#[a, 1, b] \mid [1, 0, 0, 1]$. More generally, If \mathcal{G} and \mathcal{R} are sets of signatures, and vertices of U (resp. V) are labeled by signatures from \mathcal{G} (resp. \mathcal{R}), then we also use $\#\mathcal{G} \mid \mathcal{R}$ to denote the bipartite Holant Problem. Signatures in \mathcal{G} are called *generators* and signatures in \mathcal{R} are called *recognizers*. This notation is particularly convenient when we perform holographic transformations. Throughout this paper, all $(2, 3)$ -regular bipartite graphs are arranged with generators on the degree 2 side and recognizers on the degree 3 side.

We use Arg to denote the principal value of the complex argument; i.e., $\text{Arg}(c) \in (-\pi, \pi]$ for all nonzero $c \in \mathbb{C}$.

2.1. \mathcal{F} -Gate

Any signature from \mathcal{F} is available at a vertex as part of an input graph. Instead of a single vertex, we can use graph fragments to generalize this notion. An \mathcal{F} -gate Γ is a pair (H, \mathcal{F}) , where $H = (V, E, D)$ is a graph with some dangling edges D (Figure 1 contains some examples). Other than these dangling edges, an \mathcal{F} -gate is the same as a signature grid. The role of dangling edges is similar to that of external nodes in Valiant's notion [16], however we allow more than one dangling edge for a node. In $H = (V, E, D)$ each node is assigned a function in \mathcal{F} (we do not consider “dangling” leaf nodes at the end of a dangling edge among these), E are the regular edges, and D are the dangling edges. Then we can define a function for this \mathcal{F} -gate $\Gamma = (H, \mathcal{F})$,

$$\Gamma(y_1, y_2, \dots, y_q) = \sum_{(x_1, x_2, \dots, x_p) \in \{0,1\}^p} H(x_1, x_2, \dots, x_p, y_1, y_2, \dots, y_q),$$

where $p = |E|$, $q = |D|$, $(y_1, y_2, \dots, y_q) \in \{0, 1\}^q$ denotes an assignment on the dangling edges, and $H(x_1, x_2, \dots, x_p, y_1, y_2, \dots, y_q)$ denotes the value of the signature grid on an assignment of all edges, i.e., the product of evaluations at every vertex of H , for $(x_1, x_2, \dots, x_p, y_1, y_2, \dots, y_q) \in \{0, 1\}^{p+q}$. We will also call this function the signature of the




Figure 1: Examples of binary starter, recursive, and finisher gadgets

\mathcal{F} -gate Γ . An \mathcal{F} -gate can be used in a signature grid as if it is just a single node with the same signature. We note that even for a very simple signature set \mathcal{F} , the signatures for all \mathcal{F} -gates can be quite complicated and expressive. Matchgate signatures are an example [16].

The dangling edges of an \mathcal{F} -gate are considered as input or output variables. Any m -input n -output \mathcal{F} -gate can be viewed as a 2^n by 2^m matrix M which transforms arity- m signatures into arity- n signatures (this is true even if m or n are 0). Our construction will transform symmetric signatures to symmetric signatures. This implies that there exists an equivalent $n + 1$ by $m + 1$ matrix \tilde{M} which operates directly on column vectors written in symmetric signature notation. We will henceforth identify the matrix \tilde{M} with the \mathcal{F} -gate itself. The constructions in this paper are based upon three different types of bipartite \mathcal{F} -gates which we call *starter gadgets*, *recursive gadgets*, and *finisher gadgets*. An *arity- r starter gadget* is an \mathcal{F} -gate with no input but r output edges. If an \mathcal{F} -gate has r input and r output edges then it is called an *arity- r recursive gadget*. Finally, an \mathcal{F} -gate is an *arity- r finisher gadget* if it has r input edges 1 output edge. As a matter of convention, we consider any dangling edge incident with a generator as an output edge and any dangling edge incident with a recognizer as an input edge; see Figure 1.

3. Interpolation Techniques

3.1. Binary recursive construction

In this section, we develop our new technique of higher dimensional iterations for interpolation of unary signatures.

Lemma 3.1. *Suppose $M \in \mathbb{C}^{3 \times 3}$ is a nonsingular matrix, $s \in \mathbb{C}^3$ is a nonzero vector, and for all integers $k \geq 1$, s is not a column eigenvector of M^k . Let $F_i \in \mathbb{C}^{2 \times 3}$ be three matrices, where $\text{rank}(F_i) = 2$ for $1 \leq i \leq 3$, and the intersection of the row spaces of F_i is trivial $\{0\}$. Then for every n , there exists some $F \in \{F_i : 1 \leq i \leq 3\}$, and some $S \subseteq \{FM^k s : 0 \leq k \leq n^3\}$, such that $|S| \geq n$ and vectors in S are pairwise linearly independent.*

Proof. Let $k > j \geq 0$ be integers. Then $M^k s$ and $M^j s$ are nonzero and also linearly independent, since otherwise s is an eigenvector of M^{k-j} . Let $N = [M^j s, M^k s] \in \mathbb{C}^{3 \times 2}$, then $\text{rank}(N) = 2$, and $\ker(N^T)$ is a 1-dimensional linear subspace. It follows that there exists an $F \in \{F_i : 1 \leq i \leq 3\}$ such that the row space of F does not contain $\ker(N^T)$, and hence has trivial intersection with $\ker(N^T)$. In other words, $\ker(N^T F^T) = \{0\}$. We conclude that $FN \in \mathbb{C}^{2 \times 2}$ has rank 2, and $FM^j s$ and $FM^k s$ are linearly independent.

Each F_i , where $1 \leq i \leq 3$, defines a coloring of the set $K = \{0, 1, \dots, n^3\}$ as follows: color $k \in K$ with the linear subspace spanned by $F_i M^k s$. Thus, F_i defines an equivalence relation \approx_i where $k \approx_i k'$ iff they receive the same color. Assume for a contradiction that for each F_i , where $1 \leq i \leq 3$, there are not n pairwise linearly independent vectors among $\{F_i M^k s : k \in K\}$. Then, including possibly the 0-dimensional space $\{0\}$, there can be at most n distinct colors assigned by F_i . By the pigeonhole principle, some k and k' with $0 \leq k < k' \leq n^3$ must receive the same color for all F_i , where $1 \leq i \leq 3$. This is a contradiction and we are done. \blacksquare

The next lemma says that under suitable conditions we can construct all unary signatures $[x, y]$. The method will be interpolation at a higher dimensional iteration, and finishing up with a suitable *finisher* gadget. The crucial new technique here is that when iterating at a higher dimension, we can guarantee the existence of *one* finisher gadget that succeeds on polynomially many steps, which results in overall success. Different finisher gadgets may work for different initial signatures and different input size n , but these need not be known in advance and have no impact on the final success of the reduction.

Lemma 3.2. *Suppose that the following gadgets can be built using complex-valued signatures from a finite generator set \mathcal{G} and a finite recognizer set \mathcal{R} .*

- (1) *A binary starter gadget with nonzero signature $[z_0, z_1, z_2]$.*
- (2) *A binary recursive gadget with nonsingular recurrence matrix M , for which $[z_0, z_1, z_2]^T$ is not a column eigenvector of M^k for any positive integer k .*
- (3) *Three binary finisher gadgets with rank 2 matrices $F_1, F_2, F_3 \in \mathbb{C}^{2 \times 3}$, where the intersection of the row spaces of F_1, F_2 , and F_3 is the zero vector.*

Then for any $x, y \in \mathbb{C}$, $\#\mathcal{G} \cup \{[x, y]\} \mid \mathcal{R} \leq_T \#\mathcal{G} \mid \mathcal{R}$.

Proof. The construction begins with the binary starter gadget with signature $[z_0, z_1, z_2]$, which we call N_0 . Let $\mathcal{F} = \mathcal{G} \cup \mathcal{R}$. Recursively, \mathcal{F} -gate N_{k+1} is defined to be N_k connected to the binary recursive gadget in such a way that the input edges of the binary recursive gadget are merged with the output edges of N_k . Then \mathcal{F} -gate G_k is defined to be N_k connected to one of the finisher gadgets, with the input edges of the finisher gadget merged with the output edges of N_k (see Figure 1(d)). Herein we analyze the construction with respect to a given bipartite signature grid Ω for the Holant Problem $\#\mathcal{G} \cup \{[x, y]\} \mid \mathcal{R}$, with underlying graph $G = (V, E)$. Let $Q \subseteq V$ be the set of vertices with $[x, y]$ signatures, and let $n = |Q|$. By Lemma 3.1 fix j so that at least $n + 2$ of the first $(n + 2)^3 + 1$ vectors of the form $F_j M^k [z_0, z_1, z_2]^T$ are pairwise linearly independent. We use finisher gadget F_j in the recursive construction, so that the signature of G_k is $F_j M^k [z_0, z_1, z_2]^T$, which we denote by $[X_k, Y_k]$. We note that there exists a subset S of these signatures for which each Y_k is nonzero and $|S| = n + 1$. We will argue using only the existence of S , so there is no need to algorithmically “find” such a set, and for that matter, one can try out all three finisher gadgets without any need to determine which finisher gadget is “the correct one” beforehand. If we replace every element of Q with a copy of G_k , we obtain an instance of $\#\mathcal{G} \mid \mathcal{R}$ (note that the correct bipartite signature structure is preserved), and we denote this new signature grid by Ω_k . Then

$$\text{Holant}_{\Omega_k} = \sum_{0 \leq i \leq n} c_i X_k^i Y_k^{n-i}$$

where $c_i = \sum_{\sigma \in J_i} \prod_{v \in V \setminus Q} f_v(\sigma|_{E(v)})$, J_i is the set of $\{0, 1\}$ edge assignments where the number of 0s assigned to the edges incident to the copies of G_k is i , f_v is the signature at v ,

and $E(v)$ is the set of edges incident to v . The important point is that the c_i values do not depend on X_k or Y_k . Since each signature grid Ω_k is an instance of $\#\mathcal{G} \mid \mathcal{R}$, Holant_{Ω_k} can be solved exactly using the oracle. Carrying out this process for every $k \in \{0, 1, \dots, (n+2)^3\}$, we arrive at a linear system where the c_i values are the unknowns.

$$\begin{bmatrix} \text{Holant}_{\Omega_0} \\ \text{Holant}_{\Omega_1} \\ \vdots \\ \text{Holant}_{\Omega_{(n+2)^3}} \end{bmatrix} = \begin{bmatrix} X_0^0 Y_0^n & X_0^1 Y_0^{n-1} & \cdots & X_0^n Y_0^0 \\ X_1^0 Y_1^n & X_1^1 Y_1^{n-1} & \cdots & X_1^n Y_1^0 \\ \vdots & \vdots & \ddots & \vdots \\ X_{(n+2)^3}^0 Y_{(n+2)^3}^n & X_{(n+2)^3}^1 Y_{(n+2)^3}^{n-1} & \cdots & X_{(n+2)^3}^n Y_{(n+2)^3}^0 \end{bmatrix} \begin{bmatrix} c_0 \\ c_1 \\ \vdots \\ c_n \end{bmatrix}.$$

Define $x_i = X_{k_i}$ and $y_i = Y_{k_i}$ where $S = \{k_0, k_1, \dots, k_n\}$, so that $[x_i, y_i] \in S$ for $0 \leq i \leq n$, and we have a subsystem

$$\begin{bmatrix} y_0^{-n} \cdot \text{Holant}_{\Omega_0} \\ y_1^{-n} \cdot \text{Holant}_{\Omega_1} \\ \vdots \\ y_n^{-n} \cdot \text{Holant}_{\Omega_n} \end{bmatrix} = \begin{bmatrix} x_0^0 y_0^0 & x_0^1 y_0^{-1} & \cdots & x_0^n y_0^{-n} \\ x_1^0 y_1^0 & x_1^1 y_1^{-1} & \cdots & x_1^n y_1^{-n} \\ \vdots & \vdots & \ddots & \vdots \\ x_n^0 y_n^0 & x_n^1 y_n^{-1} & \cdots & x_n^n y_n^{-n} \end{bmatrix} \begin{bmatrix} c_0 \\ c_1 \\ \vdots \\ c_n \end{bmatrix}.$$

The matrix above has entry $(x_r/y_r)^c$ at index (r, c) . Due to pairwise linear independence of $[x_r, y_r]$, x_r/y_r is pairwise distinct for each $r \in S$. Hence this is a Vandermonde system of full rank. Therefore the initial feasible linear system has full rank and we can solve it for the c_i values. With these values in hand, we can calculate $\text{Holant}_{\Omega} = \sum_{0 \leq i \leq n} c_i x^i y^{n-i}$ directly, completing the reduction. \blacksquare

The ability to simulate all unary signatures will allow us to prove $\#\text{P}$ -hardness. The next lemma says that, if \mathcal{R} contains the EQUALITY gate $=_3$, then other than on a 1-dimensional curve $ab = 1$ and an isolated point $(a, b) = (0, 0)$, the ability to simulate unary signatures gives a reduction from VERTEX COVER. Note that counting VERTEX COVER on 3-regular graphs is just $\#[0, 1, 1] \mid [1, 0, 0, 1]$. Xia et al. showed that this is $\#\text{P}$ -hard even when the input is restricted to 3-regular planar graphs [17]. We will see shortly that on the curve $ab = 1$ and at $(a, b) = (0, 0)$, the problem $\text{Hol}(a, b)$ is tractable.

Lemma 3.3. *Suppose that $(a, b) \in \mathbb{C}^2 - \{(a, b) : ab = 1\} - \{(0, 0)\}$ and let \mathcal{G} and \mathcal{R} be finite signature sets where $[a, 1, b] \in \mathcal{G}$ and $[1, 0, 0, 1] \in \mathcal{R}$. Further assume that $\#\mathcal{G} \cup \{[x_i, y_i] : 0 \leq i < m\} \mid \mathcal{R} \leq_T \#\mathcal{G} \mid \mathcal{R}$ for any $x_i, y_i \in \mathbb{C}$ and $m \in \mathbb{Z}^+$. Then $\#\mathcal{G} \cup \{[0, 1, 1]\} \mid \mathcal{R} \leq_T \#\mathcal{G} \mid \mathcal{R}$, and $\#\mathcal{G} \mid \mathcal{R}$ is $\#\text{P}$ -hard.*

Proof. Since $\#[0, 1, 1] \mid [1, 0, 0, 1]$ is $\#\text{P}$ -hard, we only need to show how to simulate the generator signature $[0, 1, 1]$. Respectively, Gadgets 1, 2, and 3 (Figure 2) can be used to simulate generator signatures $[b^{-1}, 1, 2b]$, $[0, 1, 5/(2a)]$, and $[0, 1, 1]$ in the cases where $ab = 0$, $ab = -1$, and both $ab \neq 0$ and $ab \neq -1$ (when $ab = 0$, we assume without loss of generality that $a = 0$ and $b \neq 0$). To carry this out, we set $\theta = [b, b^{-1}]$ in Gadget 1; $\theta = [1/(6a), -a/24]$ and $\gamma = [-3/a, a]$ in Gadget 2; and $\theta = (ab+1)(1-ab)^{-1}[1, -a^2]$, $\gamma = [-a^{-2}, b^{-1}(1+ab)^{-1}]$, and $\rho = (ab-1)^{-1}[-b, a]$ in Gadget 3. This results in a chain of reductions to simulate $[0, 1, 1]$ in all cases (i.e. Gadget 2 simulates a signature to be used with Gadget 1, which in turn simulates a signature to be used with Gadget 3, and Gadget 3 simulates $[0, 1, 1]$). \blacksquare

It will be helpful to have conditions that are easier to check than those in Lemma 3.2. To this end, we establish condition 2 in terms of eigenvalues, and we build general-purpose finisher gadgets to eliminate condition 3. Let M_4 , M_5 , and F be the recurrence matrices for Gadget 4, Gadget 5, and the simplest possible binary finisher gadget (each




Figure 2: Gadgets used to simulate the $[0,1,1]$ signature

built using generator signature $[a, 1, b]$ and recognizer signature $[1, 0, 0, 1]$; see Figures 3(a), 3(b), and 1(c)). Provided that $ab \neq 1$ and $a^3 \neq b^3$, it turns out that the finisher gadget sets $\{F, FM_4, FM_4^2\}$ and $\{F, FM_4, FM_5\}$ satisfy condition 3 of Lemma 3.2 when $ab \neq 0$ and $ab = 0$, respectively. Together with Lemma 3.3, these observations yield the following.

Theorem 3.4. *If the following gadgets can be built using generator $[a, 1, b]$ and recognizer $[1, 0, 0, 1]$ where $a, b \in \mathbb{C}$, $ab \neq 1$, and $a^3 \neq b^3$, then the problem $\text{Hol}(a, b)$ is $\#P$ -hard.*

- (1) A binary recursive gadget with nonsingular recurrence matrix M which has eigenvalues α and β such that $\frac{\alpha}{\beta}$ is not a root of unity.
- (2) A binary starter gadget with signature s which is not orthogonal to any row eigenvector of M . ■

3.2. Unary recursive construction

Now we consider the unary case. The following lemma arrives from [12] and is stated explicitly in [6]. It can be viewed as a unary version of Lemma 3.2 without finisher gadgets.

Lemma 3.5. *Suppose there is a unary recursive gadget with nonsingular matrix M and a unary starter gadget with nonzero signature vector s . If the ratio of the eigenvalues of M is not a root of unity and s is not a column eigenvector of M , then these gadgets can be used to interpolate all unary signatures.* ■

Surprisingly, a set of general-purpose starter gadgets can be made for this construction as long as $ab \neq 1$ and $a^3 \neq b^3$, so we refine this lemma by eliminating the starter gadget requirement. The starter gadgets are Fs , FM_4s , and FM_6s where M_6 is Gadget 6 and s is the single-vertex starter gadget (see Figures 3(c) and 1(a)).

Theorem 3.6. *Suppose there is a unary recursive gadget with nonsingular matrix M , and the ratio of the eigenvalues of M is not a root of unity. Then for any $a, b \in \mathbb{C}$ where $ab \neq 1$ and $a^3 \neq b^3$, there is a starter gadget built using generator $[a, 1, b]$ and recognizer $[1, 0, 0, 1]$ for which the resulting construction can be used to interpolate all unary signatures.* ■




Figure 3: Binary recursive gadgets

4. Complex Signatures

Now we aim to characterize $\text{Hol}(a, b)$ where $a, b \in \mathbb{C}$. The next lemma introduces the technique of algebraic symmetrization. We show that over 3-regular graphs, the Holant value is expressible as an integer polynomial $P(X, Y)$, where $X = ab$ and $Y = a^3 + b^3$. This change of variable, from (a, b) to (X, Y) , is crucial in two ways. First, it allows us to derive tractability results easily, drawing connections between problems that may appear unrelated, and the tractability of one implies the other. Second, it facilitates the proof of hardness for those (a, b) where the problem is indeed $\#P$ -hard by reducing the degree of the polynomials involved. Once this transformation is made, four binary recursive gadgets easily cover all of the $\#P$ -hard problems where X and Y are real-valued, with a straightforward symbolic computation using `CYLINDRICALDECOMPOSITION` in `Mathematica`TM. All gadget constructions in this section use $[a, 1, b]$ and $[1, 0, 0, 1]$ signatures exclusively, and we henceforth denote $X = ab$ and $Y = a^3 + b^3$ for the remainder of this paper.

Lemma 4.1. *Let G be a 3-regular graph. Then there exists a polynomial $P(\cdot, \cdot)$ with two variables and integer coefficients such that for any signature grid Ω having underlying graph G and every edge labeled $[a, 1, b]$, the Holant value is $\text{Holant}_\Omega = P(ab, a^3 + b^3)$.*

Proof. Consider any $\{0, 1\}$ vertex assignment σ with a non-zero valuation. If σ' is the complement assignment switching all 0's and 1's in σ , then for σ and σ' , we have the sum of valuations $a^i b^j + a^j b^i$ for some i and j . Here i (resp. j) is the number of edges connecting two degree 3 vertices both assigned 0 (resp. 1) by σ . We note that $a^i b^j + a^j b^i = (ab)^{\min(i,j)}(a^{|i-j|} + b^{|i-j|})$.

We prove $i \equiv j \pmod{3}$ inductively. For the all-0 assignment, this is clear since every edge contributes a factor a and the number of edges is divisible by 3 for a 3-regular graph. Now starting from any assignment σ , if we switch the assignment on one vertex from 0 to 1, it is easy to verify that it changes the valuation from $a^i b^j$ to $a^{i'} b^{j'}$, where $i - j = i' - j' + 3$. As every $\{0, 1\}$ assignment is obtainable from the all-0 assignment by a sequence of switches, the conclusion $i \equiv j \pmod{3}$ follows.

Now $a^i b^j + a^j b^i = (ab)^{\min(i,j)}(a^{3k} + b^{3k})$, for some $k \geq 0$ and a simple induction $a^{3(k+1)} + b^{3(k+1)} = (a^{3k} + b^{3k})(a^3 + b^3) - (ab)^3(a^{3(k-1)} + b^{3(k-1)})$ shows that the Holant is a polynomial $P(ab, a^3 + b^3)$ with integer coefficients. ■

Corollary 4.2. *If $ab = -1$ and $a^{12} = 1$, then $\text{Hol}(a, b)$ is in P.*

Proof. Immediate from Lemma 4.1, since the problems $\text{Hol}(1, -1)$, $\text{Hol}(-i, -i)$, and $\text{Hol}(i, i)$ are all known to be solvable in P (these fall within the families \mathcal{F}_1 , \mathcal{F}_2 , and \mathcal{F}_3 in [3]). ■

We now list all of the cases where $\text{Hol}(a, b)$ is computable in polynomial time.

Theorem 4.3. *If any of the following four conditions is true, then $\text{Hol}(a, b)$ is solvable in P:*

- (1) $ab = 1$,
- (2) $a = b = 0$,
- (3) $a^{12} = 1$ and $b = -a^{-1}$,
- (4) $a^3 = b^3$ and the input is restricted to planar graphs.

Proof. If $ab = 1$ then the signature $[a, 1, b]$ is degenerate and the Holant can be computed in polynomial time. If $a = b = 0$, a 2-coloring algorithm can be employed on the edges. If $a^{12} = 1$ and $b = -a^{-1}$ then we are done by Corollary 4.2. If we restrict the input to planar graphs and $a^3 = b^3$, holographic algorithms can be applied [4]. ■




Figure 4: Unary recursive gadgets

Our main task in this paper is to prove that all remaining problems are #P-hard. The following two lemmas provide sufficient conditions to satisfy the eigenvalue requirement of the recursive constructions.

Lemma 4.4. *If both roots of the complex polynomial $x^2 + Bx + C$ have the same norm, then $B|C| = \overline{B}C$ and $B^2\overline{C} = \overline{B^2}C$. If further $B \neq 0$ and $C \neq 0$, then $\text{Arg}(B^2) = \text{Arg}(C)$. ■*

Lemma 4.5. *If all roots of the complex polynomial $x^3 + Bx^2 + Cx + D$ have the same norm, then $C|C|^2 = \overline{B}|B|^2D$. ■*

Now we introduce a powerful new technique called *Eigenvalue Shifted Pairs*.

Definition 4.6. A pair of nonsingular square matrices M and M' is called an *Eigenvalue Shifted Pair (ESP)* if $M' = M + \delta I$ for some non-zero $\delta \in \mathbb{C}$, and M has distinct eigenvalues.

Clearly for such a pair, M' also has distinct eigenvalues. The recurrence matrices of Gadgets 10 and 11 (Figure 4) differ only by $ab-1$ along the diagonal, and form an Eigenvalue Shifted Pair for nearly all $a, b \in \mathbb{C}$. We will make significant use of such Eigenvalue Shifted Pairs, but first we state a technical lemma.

Lemma 4.7. *Suppose $\alpha, \beta, \delta \in \mathbb{C}$, $|\alpha| = |\beta|$, $\alpha \neq \beta$, $\delta \neq 0$, and $|\alpha + \delta| = |\beta + \delta|$. Then there exists $r, s \in \mathbb{R}$ such that $r\delta = \alpha + \beta$ and $s\delta^2 = \alpha\beta$. ■*

Corollary 4.8. *Let M and M' be an Eigenvalue Shifted Pair of 2 by 2 matrices. If both M and M' have eigenvalues of equal norm, then there exists $r, s \in \mathbb{R}$ such that $\text{tr}(M) = r\delta$ (possibly 0) and $\det(M) = s\delta^2$.*

Proof. Let α and β be the eigenvalues of M , so $\alpha + \delta$ and $\beta + \delta$ are the eigenvalues of M' . Suppose that $|\alpha| = |\beta|$ and $|\alpha + \delta| = |\beta + \delta|$. Then by Lemma 4.7, there exists $r, s \in \mathbb{R}$ such that $\text{tr}(M) = \alpha + \beta = r\delta$ and $\det(M) = \alpha\beta = s\delta^2$. ■

We now apply an ESP to prove that most settings of $\text{Hol}(a, b)$ are #P-hard.

Lemma 4.9. *Suppose $X \neq \pm 1$, $X^2 + X + Y \neq 0$, and $4(X - 1)^2(X + 1) \neq (Y + 2)^2$. Then either unary Gadget 10 or unary Gadget 11 has nonzero eigenvalues with distinct norm, unless X and Y are both real numbers.*

Proof. Gadgets 10 and 11 have $M_{10} = \begin{bmatrix} a^3 + 1 & a + b^2 \\ a^2 + b & b^3 + 1 \end{bmatrix}$ and $M_{11} = \begin{bmatrix} a^3 + ab & a + b^2 \\ a^2 + b & ab + b^3 \end{bmatrix}$ as their recurrence matrices, so $M_{11} = M_{10} + (X - 1)I$, and the eigenvalue shift is nonzero. Checking the determinants, $\det(M_{10}) = (X - 1)^2(X + 1) \neq 0$ and $\det(M_{11}) = (X - 1)(X^2 + X + Y) \neq 0$. Also, $\text{tr}(M_{10})^2 - 4\det(M_{10}) = (Y + 2)^2 - 4(X - 1)^2(X + 1) \neq 0$, so the eigenvalues of M_{10} are distinct. Therefore by Corollary 4.8, either M_{10} or M_{11} has nonzero eigenvalues of distinct norm unless $\text{tr}(M_{10}) = r(X - 1)$ and $\det(M_{10}) = s(X - 1)^2$ for some $r, s \in \mathbb{R}$. Then we would have $(X - 1)^2(X + 1) = s(X - 1)^2$ so $X = s - 1 \in \mathbb{R}$ and $Y + 2 = r(X - 1)$ so $Y = r(X - 1) - 2 \in \mathbb{R}$. ■

Now we will deal with the following exceptional cases from Lemma 4.9 ($X = 1$ is tractable by Theorem 4.3).

- 0. $X \in \mathbb{R}$ and $Y \in \mathbb{R}$
- 1. $X^2 + X + Y = 0$
- 2. $X = -1$
- 3. $4(X - 1)^2(X + 1) = (Y + 2)^2$

The case where X and Y are both real is dealt with using the tools developed in Section 3, and some symbolic computation. This includes the case where a and b are both real as a subcase. When a and b are both real, a dichotomy theorem for the complexity of $\text{Hol}(a, b)$ has been proved in [6] with a significant effort. With the new tools developed, we offer a simpler proof. This also covers some cases where a or b is complex. Working with real-valued X and Y is a significant advantage, since the failure condition given by Lemma 4.5 is simplified by the disappearance of norms and conjugates. This brings the problem of proving #P-hardness within reach of symbolic computation via cylindrical decomposition. We apply Theorem 3.4 to Gadgets 4, 7, 8, and 9 (Figure 3) together with a starter gadget (Figure 1(a)) to prove that these problems are #P-hard. Conditions 1 and 2 of Theorem 3.4 are encoded directly into a query for CYLINDRICALDECOMPOSITION in Mathematica™.

Theorem 4.10. *Suppose $a, b \in \mathbb{C}$, $X, Y \in \mathbb{R}$, $ab \neq 1$, $a^3 \neq b^3$, and it is not the case that $a^6 = 1$ and $ab = -1$. Then the problem $\text{Hol}(a, b)$ is #P-hard.* ■

Now we can assume that $X \notin \mathbb{R}$ or $Y \notin \mathbb{R}$, and we deal with the remaining three conditions. Note that if $X^2 + X + Y = 0$ then $X \in \mathbb{R}$ implies $Y \in \mathbb{R}$. So in the following lemma, the assumption that X and Y are not both real numbers amounts to $X \notin \mathbb{R}$.

Lemma 4.11. *If $X^2 + X + Y = 0$ and $X \notin \mathbb{R}$ then the recurrence matrix of unary Gadget 12 has nonzero eigenvalues with distinct norm.*

Proof. Let M_{12} be the recurrence matrix for unary Gadget 12. Then $\det(M_{12}) = X^6 - 6X^5 - X^4Y + 16X^4 + 11X^3Y - 10X^3 + 5X^2Y^2 - 7X^2Y - X^2 + XY^3 - 4XY^2 - 3XY - Y^3 - Y^2$. Amazingly, with the condition $X^2 + X + Y = 0$, this polynomial factors into $-X^2(X - 1)^5$. Similarly, the trace, which is $-2X^3 + 6X^2 + 3XY + Y^2 + Y$, also factors into $X(X - 1)^3$. Since $\det(M_{12}) \neq 0$, $\text{tr}(M_{12}) \neq 0$, and $(1 - X)\det(M_{12}) = \text{tr}(M_{12})^2$, we know $\text{Arg}(\det(M_{12})) \neq \text{Arg}(\text{tr}(M_{12})^2)$ and conclude by Lemma 4.4 that the eigenvalues of M_{12} (which are nonzero) have distinct norm. ■

Similarly, Gadgets 11 and 13 can be used to deal with the $X = -1$ condition. The $4(X - 1)^2(X + 1) = (Y + 2)^2$ condition can be dealt with using Gadgets 13 and 14 (an ESP) in addition to Gadgets 15 and 16. These gadgets, together with Lemma 3.3 and Theorem 3.6, give the following theorem (note that $X = -1$ and $Y = \pm 2i$ if and only if $a^3 = \pm i$ and $b = -a^{-1}$; any such setting of a and b is tractable by Theorem 4.3).

Theorem 4.12. *Suppose $a, b \in \mathbb{C}$ such that X and Y are not both real, $X \neq 1$, $a^3 \neq b^3$, and either $X \neq -1$ or $Y \neq \pm 2i$. Then the problem $\text{Hol}(a, b)$ is #P-hard.* ■

Recall VERTEX COVER is #P-hard on 3-regular planar graphs, and note that all gadgets discussed are planar (in the case of Gadget 8, each iteration can be redrawn in a planar way by “going around” the previous iterations; see Figure 1(d)). Thus, all of the hardness results proved so far still apply when the input graphs are restricted to planar graphs. There are, however, a few cases where the problem is #P-hard in general, yet is polynomial time computable when restricted to planar graphs. The relevant interpolation results can be

obtained with Gadget 4 and holographic reductions, using a technique demonstrated in [5]. Given this, we have the following result.

Theorem 4.13. *The problem $\text{Hol}(a, b)$ is $\#P$ -hard for all $a, b \in \mathbb{C}$ except in the following cases, for which the problem is in P .*

- (1) $ab = 1$
- (2) $a = b = 0$
- (3) $a^{12} = 1$ and $b = -a^{-1}$

If we restrict the input to planar graphs, then these three categories are tractable in P , as well as a fourth category $a^3 = b^3$, and the problem remains $\#P$ -hard in all other cases. ■

A simple coordinate change from (a, b) to $(X, (\frac{Y}{2})^2)$ translates this into Theorem 1.1.

References

- [1] L. Blum, F. Cucker, M. Shub, and S. Smale. *Complexity and real computation*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1998.
- [2] A. Bulatov and M. Grohe. The complexity of partition functions. *Theoretical Computer Science*, 348(2-3):148–186, 2005.
- [3] J-Y. Cai, X. Chen, and P. Lu. Graph homomorphisms with complex values: a dichotomy theorem. *CoRR*, abs/0903.4728, 2009.
- [4] J-Y. Cai and P. Lu. Holographic algorithms: from art to science. In *STOC '07: Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 401–410, 2007.
- [5] J-Y. Cai, P. Lu, and M. Xia. Holographic algorithms by Fibonacci gates and holographic reductions for hardness. In *FOCS '08: Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 644–653, 2008.
- [6] J-Y. Cai, P. Lu, and M. Xia. A computational proof of complexity of some restricted counting problems. In *TAMC '09: Proceedings of Theory and Applications of Models of Computation, 6th Annual Conference*, LNCS 5532, pages 138–149, 2009.
- [7] M. Dyer, L.A. Goldberg, and M. Paterson. On counting homomorphisms to directed acyclic graphs. *Journal of the ACM*, 54(6), 2007.
- [8] M. Dyer and C. Greenhill. The complexity of counting graph homomorphisms (extended abstract). In *SODA '00: Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 246–255, 2000.
- [9] L.A. Goldberg, M. Grohe, M. Jerrum, and M. Thurley. A complexity dichotomy for partition functions with mixed signs. *CoRR*, abs/0804.1932, 2008.
- [10] P. Hell and J. Nešetřil. On the complexity of H -coloring. *Journal of Combinatorial Theory, Series B*, 48(1):92–110, 1990.
- [11] M. Kowalczyk and J-Y. Cai. Holant Problems for Regular Graphs with Complex Edge Functions. *CoRR*, abs/1001.0464, 2010.
- [12] S. Vadhan. The complexity of counting in sparse, regular, and planar graphs. *SIAM Journal on Computing*, 31(2):398–427, 2001.
- [13] L. Valiant. The complexity of enumeration and reliability problems. *SIAM Journal on Computing*, 8(3):410–421, 1979.
- [14] L. Valiant. The complexity of computing the permanent. *Theoretical Computer Science* 8:189–201, 1979.
- [15] L. Valiant. Holographic algorithms. *SIAM Journal on Computing*, 37(5):1565–1594, 2008.
- [16] L. Valiant. Quantum circuits that can be simulated classically in polynomial time. *SIAM Journal on Computing*, 31(4): 1229–1254, 2002.
- [17] M. Xia, P. Zhang, and W. Zhao. Computational complexity of counting problems on 3-regular planar graphs. *Theoretical Computer Science*, 384(1):111–125, 2007.