Holant Problems for Regular Graphs with Complex Edge Functions

Michael Kowalczyk* Jin-Yi Cai†

Abstract

We prove a complexity dichotomy theorem for the following class of Holant Problems. Given a 3-regular graph $G = (V, E)$, compute

$$\text{Holant}(G) = \sum_{\sigma:V \to \{0,1\}} \prod_{\{u,v\} \in E} g(\{\sigma(u), \sigma(v)\}),$$

where the (symmetric) edge function g is arbitrary complex valued. 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. These theorems can be extended to k-regular graphs. With *holographic reductions* the classification theorem also applies to problems beyond the basic model.

Keywords: Computational complexity

1 Introduction

In this paper we consider the following subclass of Holant Problems [4, 5]. 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 $\text{Holant}(G) = \sum_{\sigma:V \to \{0,1\}} \prod_{\{u,v\} \in E} g(\{\sigma(u), \sigma(v)\})$.

Holant Problems are a natural class of counting problems. As introduced in [4, 5], the general Holant Problem framework can encode all counting Constrained 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 [7, 2, 6, 8, 3, 9] 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 [7] and [6] respectively. A dichotomy theorem for any symmetric matrix H with non-negative real entries is proved in [2]. Goldberg et. al. [8] 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].

*Northern Michigan University mkowalcz@nmu.edu
†University of Wisconsin-Madison jyc@cs.wisc.edu. Supported by NSF CCF-0830488 and CCF-0914969.

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 $x = g(00), y = g(01) = g(10)$ and $z = g(11)$. 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 it and assume $y = 1$. Our main theorem is as follows:

Theorem 1.1. *Suppose $a, b \in \mathbb{C}$, and let $X = ab$, $Q = (\frac{a^3+b^3}{2})^2$. 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. $X = Q = 0$
3. $X = -1$ *and* $Q = 0$
4. $X = -1$ *and* $Q = -1$

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

These results can be extended to k-regular graphs. One can also use holographic reductions [13] 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 [11, 12] introduced the powerful technique of *interpolation*, which was further developed by many others. In [4] a dichotomy theorem is proved for the case when g is a 0-1 valued Boolean function. The technique from [4] 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 are powerful enough to show that they succeed in each case. Nonetheless this involves a case-by-case verification. In [5] 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 [4] 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

[1] Technically, 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.

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 λI where $\lambda \neq 0$. 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, the Dichotomy Theorem for 3-regular graphs. Some proof details are presented in the Appendix.

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 $\xi : E \to \{0, 1\}$; 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 compute[2]

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

Suppose G is a bipartite graph (U, V, E) such that each $u \in U$ on the LHS 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 $\phi : V \to \{0, 1\}$. In fact, we can merge 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 $E' = \{e_u \mid u \in U\}$. For an edge-labeled graph (V, E') where $e \in E'$ has label g_e, $\text{Holant}_\Omega = \sum_{\phi:V \to \{0,1\}} \prod_{e=(v,w) \in E'} g_e(\phi(v), \phi(w))$. If each g_e is the same function g (but assignments $\phi : V \to [q]$ 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.

[2]The term Holant was first introduced by Valiant in [13] to denote a related exponential sum.

A symmetric function $g : \{0,1\}^k \to \mathbb{C}$ can be denoted as $[g_0, g_1, \ldots, g_k]$, 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 $f = [f_0, f_1, f_2, f_3]$, then we also use $\#[g_0, g_1, g_2] \mid [f_0, f_1, f_2, f_3]$ 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 $\mathrm{Hol}(a,b) = \#[a,1,b] \mid [1,0,0,1]$. 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.

We use Arg to denote the principal value of the complex argument, i.e. $\mathrm{Arg}(c) \in (-\pi, \pi]$ 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 [14], 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, denoted as $1, 2, \ldots, m$, and D are the dangling edges, denoted as $m+1, m+2, \ldots, m+n$. Then we can define a function for this $\mathcal{F}$-gate $\Gamma = (H, \mathcal{F})$,

$$\Gamma(y_1, y_2, \ldots, y_n) = \sum_{x_1 x_2 \cdots x_m \in \{0,1\}^m} H(x_1 x_2 \cdots x_m y_1 y_2 \cdots y_n),$$

where $(y_1, y_2, \ldots, y_n) \in \{0,1\}^n$ denotes an assignment on the dangling edges and $H(x_1 x_2 \cdots x_m y_1 y_2 \cdots y_n)$ 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 $(x_1, x_2, \ldots, x_m, y_1, y_2, \ldots, y_n) \in \{0,1\}^{m+n}$. We will also call this function the signature of the $\mathcal{F}$-gate Γ. An $\mathcal{F}$-gate can be used in a signature grid as if it is just a single node with the particular 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 [14].

In the language of holographic algorithms, the signatures of an $\mathcal{F}$-gate used on one side of the bipartite graph are called *generators* and on the other side are called *recognizers*. Throughout this paper, we will consider the side with degree 2 vertices as the generators and the other side as the recognizers. 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 $\widetilde{M}$ which operates directly on column vectors written in symmetric signature notation. We will henceforth identify the matrix $\widetilde{M}$ 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

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

3.1 Binary recursive construction

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

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

Each F_i, where $1 \leq i \leq 3$, defines a coloring of the set $K = \{0, 1, \ldots, n^3\}$ as follows: color $k \in K$ with the linear subspace spanned by F_iA^kv. 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 $\{F_iA^kv : k \in K\}$. 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. □

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 by 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 A, for which $[z_0, z_1, z_2]^{\mathrm{T}}$ is not a column eigenvector of A^k for any positive integer k.*
3. *Three 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_{s+1} is defined to be N_s 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_s. Then $\mathcal{F}$-gate G_s is defined to be N_s connected to one of the finisher gadgets, with the input edges of the finisher gadget merged with the output edges of N_s (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 $F_j A^s [z_0\ z_1\ z_2]^{\mathrm{T}}$ are pairwise linearly independent. We use finisher gadget F_j in the recursive construction, so that the signature of G_s is $F_j A^s [z_0\ z_1\ z_2]^{\mathrm{T}}$, which we denote by $[X_s, Y_s]$. We note that there exists a subset S of these signatures for which each Y_s 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_s, 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_s. Then

$$\mathrm{Holant}_{\Omega_s} = \sum_{0 \le i \le n} c_i X_s^i Y_s^{n-i}$$

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_s 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_s or Y_s. Since each signature grid Ω_s is an instance of $\#\mathcal{G} \mid \mathcal{R}$, $\mathrm{Holant}_{\Omega_s}$ can be solved exactly using the oracle. Carrying out this process for every $s \in \{0,1,\ldots,(n+2)^3\}$, we arrive at a linear system where the c_i values are the unknowns.

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

Define $x_i = X_{s_i}$ and $y_i = Y_{s_i}$ where $S = \{s_0, s_1, \ldots, s_n\}$, so that $[x_i, y_i] \in S$ for $0 \le i \le n$, and we have a subsystem

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

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 $\mathrm{Holant}_\Omega = \sum_{0 \le i \le n} c_i x^i y^{n-i}$ directly, completing the reduction. □

The ability to simulate all unary signatures will allow us to prove #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 #P-hard even when the input is restricted to 3-regular planar graphs [15]. We will see shortly that on the curve $ab=1$ and at $(a,b)=(0,0)$, the problem $\mathrm{Hol}(a,b)$ 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* $\#\mathcal{G} \cup \{[x_i, y_i] : 0 \le i < m\} \mid \mathcal{R} \le_T \#\mathcal{G} \mid \mathcal{R}$ *for any* $x_i, y_i \in \mathbb{C}$ *and* $m \in \mathbb{Z}^+$. *Then* $\#\mathcal{G} \cup \{[0,1,1]\} \mid \mathcal{R} \le_T \#\mathcal{G} \mid \mathcal{R}$, *and* $\#\mathcal{G} \mid \mathcal{R}$ *is* #P*-hard.*

Proof. Assume $ab \neq 1$ and $(a,b) \neq (0,0)$. Since $\mathrm{Hol}(0,1)$ (which is the same as $\#[0,1,1] \mid [1,0,0,1]$, or counting vertex covers on 3-regular graphs) is #P-hard, we only need to show how to simulate the generator signature $[0,1,1]$. We split this into three cases, and use a chain of three reductions.

1. $ab \neq 0$ and $ab \neq -1$
2. $ab = 0$
3. $ab = -1$

If $ab \neq 0$ and $ab \neq -1$, then we use Gadget 3 (Figure 2(c)), and we set its signatures to be $\alpha = [a,1,b]$, $\beta = [(ab+1)/(1-ab), -a^2(ab+1)/(1-ab)]$, $\gamma = [-a^{-2}, b^{-1}(1+ab)^{-1}]$, and $\delta = [-b/(ab-1), a/(ab-1)]$. Calculating the resulting signature of Gadget 3, we find that it is $[0,1,1]$ as desired.

If $ab = 0$ then assume without loss of generality that $a = 0$ and $b \neq 0$. This time we use Gadget 1, setting $\alpha = [a,1,b]$ and $\beta = [b, b^{-1}]$. Then Gadget 1 simulates a $[b^{-1}, 1, 2b]$ generator signature, but since this signature fits the criteria of case 1 above, we are done by reduction from that case.

Similarly, if $ab = -1$, then Gadget 2 exhibits a generator signature of the form $[0, 1, 5/(2a)]$ under the signatures $\alpha = [a,1,b]$, $\beta = [1/(6a), -a/24]$, and $\gamma = [-3/a, a]$. Since $5/(2a)$ is nonzero, we are done by reduction from case 2. □

Theorem 3.1. *Suppose that 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* $a^3 \neq b^3$.

1. *A binary recursive gadget with nonsingular recurrence matrix* A *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* A.

Then the problem $\mathrm{Hol}(a,b)$ *is* #P*-hard.*

Proof. First we show how to build general-purpose binary finisher gadgets for the main construction using the assumed generator and recognizer, starting first with the case where $ab \neq 0$. Using the simplest possible choice for a finisher gadget F (Figure 1(c)), we get $F = \begin{bmatrix} a & 0 & 1 \\ 1 & 0 & b \end{bmatrix}$. Let M_4 be the recurrence matrix for binary recursive Gadget 4 (Figure 3(a)), and we build two more finisher gadgets F' and F'' using Gadget 4 so that $F' = FM_4$ and $F'' = FM_4^2$. Since F and M_4 both have full rank (note $\det(M_4) = ab(ab-1)^3$), it follows that F' and F'' also have full rank. Now we will show that the row spaces of F, F' and F'' have trivial intersection, and it suffices to verify that the cross products of the row vectors of F, F', and F'' are linearly independent. (To see this, note that the cross product is orthogonal to a vector if and only if that vector is in the row space. If all three cross products are linearly independent, the matrix of cross products has full rank and trivial kernel, and the kernel is precisely the intersection of the row spaces). The cross product of the row vectors of F, F', and F'' are $[0, 1-ab, 0]$, $(ab-1)^2[2b^2, -ab(1+ab), 2a^2]$, and $(ab-1)^3[2b(a^2+ab^2+a^2b^3+b^4), -ab(2a^3+ab+2a^2b^2+2b^3+a^3b^3), 2a(a^4+a^2b+b^2+a^3b^2)]$ respectively. Then to see that these 3 vectors are linearly independent, it suffices to verify that the submatrix $\begin{bmatrix} 2b^2 & 2a^2 \\ 2b(a^2+ab^2+a^2b^3+b^4) & 2a(a^4+a^2b+b^2+a^3b^2) \end{bmatrix}$ is nonsingular. Since $a \neq 0$ and $b \neq 0$, we can just verify $\det\left(\begin{bmatrix} b & a \\ (a^2+ab^2+a^2b^3+b^4) & (a^4+a^2b+b^2+a^3b^2) \end{bmatrix}\right) = (ab-1)(a^3-b^3) \neq 0$, and the matrix is nonsingular.

If $ab = 0$, assume without loss of generality that $a \neq 0$ and $b = 0$. Let M_5 be the recurrence matrix for binary recursive Gadget 5. Composing F with M_5, we get a finisher gadget with matrix FM_5, for which the cross product of the row vectors is $[-2a, 1+2a^3, -2a^2(1+a)(1-a+a^2)]$. The cross product of the rows of F and F' in this case are $[0,1,0]$ and $[0,0,2a^2]$ respectively. Then the matrix of cross products is clearly nonsingular, and we conclude that for any $a, b \in \mathbb{C}$, we have 3 finisher gadgets satisfying Lemma 3.2 item 3 unless $ab = 1$ or $a^3 = b^3$.

Now we want to show that s is not an eigenvector of A^k for any positive integer k (note that s is nonzero by assumption). Writing out the Jordan Normal Form for A, we have $A^k s = T^{-1}D^kTs$, where D^k has the form $\begin{bmatrix} \alpha^k & 0 & 0 \\ 0 & \beta^k & 0 \\ 0 & * & * \end{bmatrix}$. Let $t = Ts$ and write $t = [c\ d\ e]^{\mathrm{T}}$. By hypothesis, s is not orthogonal to the first two rows of T, thus $c, d \neq 0$. If s were an eigenvector of A^k for some positive integer k, then $T^{-1}D^kTs = A^ks = \lambda s$ for some nonzero complex value λ, and $D^kt = T\lambda s = \lambda t$. But then $c\alpha^k = \lambda c$ and $d\beta^k = \lambda d$, which means $\frac{\alpha^k}{\beta^k} = 1$, contradicting the fact that $\frac{\alpha}{\beta}$ is not a root of unity.

We have now met all the criteria for Lemma 3.2, so the reduction $\#\mathcal{S} \cup \{[a,1,b],[x,y]\} \mid [1,0,0,1] \leq_T \#\mathcal{S} \cup \{[a,1,b]\} \mid [1,0,0,1]$ holds for any $x, y \in \mathbb{C}$, and by Lemma 3.3 the problem $\mathrm{Hol}(a,b)$ is #P-hard. □

3.2 Unary recursive construction

Now we state a similar theorem for the unary case. The following lemma arrives from [10] and is stated explicitly in [5]. It can be viewed as a unary version of Lemma 3.2 without finisher gadgets.

Lemma 3.4. *Suppose there is a recursive unary gadget with nonsingular matrix A and a unary starter gadget with nonzero column vector s. If the ratio of the eigenvalues of A is not a root of unity and s is not a column eigenvector of A, then these gadgets can be used to interpolate all unary signatures.*

Surprisingly, it turns out that a set of "general purpose" starter gadgets can be made for this construction, so we refine this lemma by eliminating the starter gadget requirement. The proof of the following lemma is listed in the appendix.

Theorem 3.2. *Suppose there is a recursive unary gadget with nonsingular matrix A, built using generator $[a,1,b]$ and recognizer $[1,0,0,1]$. If the ratio of the eigenvalues of A is not a root of unity then this gadget can be used to interpolate all unary signatures, unless $ab = 1$ or $a^3 = b^3$.*

4 Complex Signatures

Now we aim to characterize $\mathrm{Hol}(a,b)$ 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. Second it facilitates the proof of hardness for those (a,b) where the problem is indeed #P-hard. Viewing the Holant in this way reduces the degree of the polynomials involved and draws connections between problems that may appear unrelated, and the tractability of one implies the other. Once this transformation is made, Gadgets 4, 6, and 7 (Figure 3) easily cover all of the #P-hard problems where X and Y are real valued, with a straightforward symbolic computation using CylindricalDecomposition in Mathematica$^{\mathrm{TM}}$.

Lemma 4.1. *Let G be a 3-regular graph. Then there exists a polynomial $P(\cdot,\cdot)$ with integer coefficients in two variables, 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 $3 \mid i-j$ 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 $3 \mid i-j$ 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 4.1. *If $ab=-1$ and $a^{12}=1$, then* $\text{Hol}(a,b)$ *is in* P.

Proof. The problems $\text{Hol}(1,-1)$, $\text{Hol}(-i,-i)$, and $\text{Hol}(i,i)$ are all solvable in P (these are listed as the families $\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3$ in [5, 3].) and the value of a^3+b^3 for these problems is 0, $2i$, and $-2i$ respectively. Moreover, a^3+b^3 takes on one of these three values for all 12 problems under consideration. Since the value of any 3-regular signature grid is completely determined by a^3+b^3, ab, and the polynomial $P(\cdot,\cdot)$ (which in turn depends only on the underlying graph G), each of these 12 problems is equivalent to either $\text{Hol}(1,-1)$, $\text{Hol}(-i,-i)$, or $\text{Hol}(i,i)$, and is thus solvable in P. □

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

Theorem 4.1. *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 4.1. If we restrict the input to planar graphs and $a^3=b^3$, holographic algorithms can be applied [4]. □

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. Proofs are listed in the Appendix.

Lemma 4.2. *If both roots of a polynomial x^2+Bx+C with $B, C \in \mathbb{C}$ have the same norm, then $B|C| = \overline{B}C$. If further $B \neq 0$ and $C \neq 0$, then* $\text{Arg}(B^2) = \text{Arg}(C)$.

Lemma 4.3. *If all roots of the complex polynomial x^3+Bx^2+Cx+D have the same norm, then $C|C|^2 = \overline{B}|B|^2 D$.*

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

Definition 4.1. *A pair of nonsingular square matrices M and M' is called an Eigenvalue Shifted Pair if $M' = M + \delta I$ for some non-zero $\delta \in \mathbf{C}$, and M has distinct eigenvalues.*

Clearly for such a pair, M' also has distinct eigenvalues.

The recurrence matrices of Gadgets 9 and 10 (Figure 4) are an example of such an Eigenvalue Shifted Pair. Their recurrence matrices differ by $ab - 1$ along the diagonal and are identical elsewhere. We will make significant use of such *Eigenvalue Shifted Pairs*. We start with a technical lemma.

Lemma 4.4. *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 $s\delta^2 = \alpha\beta$.*

Proof. After a rotation in the complex plane, we can assume $\alpha = \overline{\beta}$, and then since $\alpha + \beta, \alpha\beta \in \mathbb{R}$ we just need to prove $\delta \in \mathbb{R}$. Then $(\alpha + \delta)\overline{(\alpha + \delta)} = |\alpha + \delta|^2 = |\beta + \delta|^2 = (\beta + \delta)\overline{(\beta + \delta)} = (\overline{\alpha} + \delta)(\alpha + \overline{\delta})$ and we distribute to get $\alpha\overline{\alpha} + \delta\overline{\delta} + \alpha\overline{\delta} + \overline{\alpha}\delta = \alpha\overline{\alpha} + \delta\overline{\delta} + \overline{\alpha}\overline{\delta} + \alpha\delta$. Canceling repeated terms and factoring, we have $(\overline{\alpha} - \alpha)(\overline{\delta} - \delta) = 0$, and since $\alpha \neq \beta = \overline{\alpha}$ we know $\overline{\delta} = \delta$ therefore $\delta \in \mathbb{R}$. □

Corollary 4.2. *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 $\mathrm{tr}(M) = r\delta$ (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 4.4, there exists $r, s \in \mathbb{R}$ such that $\mathrm{tr}(M) = \alpha + \beta = r\delta$ and $\det(M) = \alpha\beta = s\delta^2$. □

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

Lemma 4.5. *Suppose $X \neq \pm 1$, $4(X-1)^2(X+1) \neq (Y+2)^2$, $Y + 2 \neq 0$, and $X^2 + X + Y \neq 0$. Then either unary Gadget 9 or unary Gadget 10 has nonzero eigenvalues with distinct norm, unless X and Y are both real numbers.*

Proof. The recurrence matrices for unary Gadgets 9 and 10 are

$$M_9 = \begin{bmatrix} 1 + a^3 & a + b^2 \\ a^2 + b & 1 + b^3 \end{bmatrix}$$
$$M_{10} = \begin{bmatrix} a^3 + ab & a + b^2 \\ a^2 + b & ab + b^3 \end{bmatrix}$$

so $M_{10} = M_9 + (X - 1)I$, and the eigenvalue shift is nonzero. Checking the determinants, $\det(M_9) = (X-1)^2(X+1) \neq 0$ and $\det(M_{10}) = (X-1)(X^2 + X + Y) \neq 0$. Also, $\mathrm{tr}(M_9)^2 - 4\det(M_9) = (Y+2)^2 - 4(X-1)^2(X+1) \neq 0$, so the eigenvalues of M_9 are distinct. Therefore by Corollary 4.2, either M_9 or M_{10} has nonzero eigenvalues of distinct norm unless $\mathrm{tr}(M_9) = r(X-1)$ and $\det(M_9) = s(X-1)^2$ 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 4.5 ($X = 1$ is tractable by Theorem 4.1).

1. $X \in \mathbb{R}$ and $Y \in \mathbb{R}$
2. $Y = -2$

3. $X^2 + X + Y = 0$

4. $X = -1$

5. $4(X-1)^2(X+1) = (Y+2)^2$

The case where X and Y are both real will be dealt with using the tools developed in Section 3, and some symbolic computation. Where both X and Y are real 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 [5] 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 4.3 is simplified by the disappearance of conjugates. This brings the problem of proving #P-hardness for all relevant problems within reach of symbolic computation via cylindrical decomposition. We apply Theorem 3.1 together with a careful selection of three binary recursive gadgets and a starter gadget to prove that these problems are #P-hard. Conditions 1 and 2 of Theorem 3.1 are encoded directly into a query for CylindricalDecomposition in Mathematica$^{\text{TM}}$. The details of the gadgets involved and how the query is formulated are both in the appendix.

Theorem 4.2. *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 4 conditions.

Lemma 4.6. *If $Y = -2$ and $X \notin \mathbb{R}$ then the recurrence matrix of unary Gadget 11 has nonzero eigenvalues with distinct norm.*

Proof. Let M_{11} be the recurrence matrix for unary Gadget 11. Then $\det(M_{11}) = (X^2 + X + Y)(X - 1)(2X+Y)^2 = 4(X-1)^4(X+2)$ and $\text{tr}(M_{11}) = (2X+Y)^2 = 4(X-1)^2$, and note these are both nonzero. If $\text{Arg}(\det(M_{11})) = \text{Arg}(\text{tr}(M_{11})^2)$, then there exists $r \in \mathbb{R}$ such that $4(X-1)^4(X+2) = 16r(X-1)^4$, so $X + 2 = 4r$, but since $X \notin \mathbb{R}$ this cannot be, so we are done by Lemma 4.2. □

The conditions $X^2 + X + Y = 0$ and $X = -1$ can be similarly dealt with through the use of carefully chosen unary recursive gadgets. The $4(X-1)^2(X+1) = (Y+2)^2$ condition can be dealt with using another application of ESP along with some other gadgets. The details are left to the appendix. The net effect is summarized in the following lemma.

Lemma 4.7. *Assume $a, b \in \mathbb{C}$ such that X and Y are not both real, $X \neq 1$, and either $X \neq -1$ or $Y \neq \pm 2i$. Then at least one of the unary recursive Gadgets 9 through 17 has a recurrence matrix with nonzero eigenvalues of distinct norm.*

Note that $X = -1$ and $Y = \pm 2i$ if and only if $a^3 = \pm i$ and $b = -1/a$; any such setting of a and b is tractable by Theorem 4.1. We have the following summary theorem.

Theorem 4.3. *Suppose $X, Y \in \mathbb{C}$ but 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.*

Proof. Immediate from Lemma 4.7 and Theorem 3.2. □

Recall Vertex Cover is #P-hard on 3-regular planar graphs, and note that all gadgets discussed are planar (in the case of Gadget 7, 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 details on planar graphs are listed in the Appendix. Given this, we have the following result.

Theorem 4.4. *The problem* 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.*

It is straightforward to change coordinates from (a, b) to $(X, (\frac{Y}{2})^2)$, and this results in Theorem 1.1.

References

[1] Lenore Blum, Felipe Cucker, Michael Shub, and Steve Smale. *Complexity and real computation.* Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1998.

[2] Andrei A. Bulatov and Martin Grohe. The complexity of partition functions. *Theor. Comput. Sci.*, 348(2-3):148–186, 2005.

[3] Jin-Yi Cai, Xi Chen, and Pinyan Lu. Graph homomorphisms with complex values: A dichotomy theorem. *manuscript*, 2009.

[4] Jin-Yi Cai, Pinyan Lu, and Mingji 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.

[5] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. A Computational Proof of Complexity of Some Restricted Counting Problems TAMC 2009: 138-149.

[6] Martin E. Dyer, Leslie Ann Goldberg, and Mike Paterson. On counting homomorphisms to directed acyclic graphs. *J. ACM*, 54(6), 2007.

[7] Martin E. Dyer and Catherine S. Greenhill. The complexity of counting graph homomorphisms (extended abstract). In *SODA*, pages 246–255, 2000.

[8] Leslie Ann Goldberg, Martin Grohe, Mark Jerrum, and Marc Thurley. A complexity dichotomy for partition functions with mixed signs. *CoRR*, abs/0804.1932, 2008.

[9] P. Hell and J. Nešetřil. On the complexity of h-coloring. *Journal of Combinatorial Theory, Series B*, 48(1):92–110, 1990.

[10] Salil P. Vadhan. The complexity of counting in sparse, regular, and planar graphs. *SIAM J. Comput.*, 31(2):398–427, 2001.

[11] Leslie G. Valiant. The complexity of enumeration and reliability problems. *SIAM J. Comput.*, 8(3):410–421, 1979.

[12] L. G. Valiant: The Complexity of Computing the Permanent. Theor. Comput. Sci. 8: 189-201 (1979)

[13] Leslie G. Valiant. Holographic algorithms (extended abstract). In *FOCS '04: Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pages 306–315,2004. SIAM J. Comput. 37(5): 1565-1594 (2008).

[14] Leslie G. Valiant: Quantum Circuits That Can Be Simulated Classically in Polynomial Time. *SIAM J. Comput.*, 31(4): 1229-1254 (2002).

[15] Mingji Xia, Peng Zhang, and Wenbo Zhao. Computational complexity of counting problems on 3-regular planar graphs. *Theor. Comput. Sci.*, 384(1):111–125, 2007.

Appendix

Proof of Theorem 3.2

Proof. Let F be the matrix of the binary finisher gadget in Figure 1(c), let M_4 be the matrix for Gadget 4, let M_8 be the matrix for binary recursive Gadget 8, and let S be the matrix for the single-vertex binary starting gadget (Figure 1(a)). Using Mathematica$^{\text{TM}}$, we verify that the matrices $[FM_4S, FS]$, $[FM_8S, FS]$, and $[FM_4S, FM_8S]$ are all nonsingular provided that $ab \neq 1$ and $a^3 \neq b^3$, so the vectors FS, FM_4S, and FM_8S are pairwise linearly independent. Since A has at most two eigenvectors at least one element of $\{FAS, FM_8S, FS\}$ is not a column eigenvector of A and can be used as a starter gadget for A. The result then follows from Lemma 3.4. □

Proof of Lemma 4.2

Proof. If the roots have equal norm, then for some $r \in \mathbb{R}$ and $a, b \in \mathbb{C}$ we can write $x^2 + Bx + C = (x - ra)(x - rb)$, where $|a| = |b| = 1$, so $B|C| = -r(a + b)r^2 = -r(a^{-1} + b^{-1})r^2ab = \overline{B}C$. Equivalently, $B^2|C| = |B^2|C$, and if B and C are nonzero, we have $\frac{B^2}{|B^2|} = \frac{C}{|C|}$, that is, $\text{Arg}(B^2) = \text{Arg}(C)$. □

Proof of Lemma 4.3

Proof. If the roots have equal norm, then for some $r \in \mathbb{R}$ and $a, b, c \in \mathbb{C}$ we can write $x^3+Bx^2+Cx+D = (x - ra)(x - rb)(x - rc)$, where $|a| = |b| = |c| = 1$, so $B = -r(a + b + c)$, $C = r^2(ab + bc + ca)$, and $D = -r^3abc$. Then

$$C|C|^2 = r^2(ab + bc + ca)r^4|ab + bc + ca|^2 = r(\overline{a + b + c})r^2|a + b + c|^2r^3abc = \overline{B}|B|^2D,$$

where we used the fact that $|ab+bc+ca| = |ab+bc+ca| \cdot |a^{-1}b^{-1}c^{-1}| = |a^{-1}+b^{-1}+c^{-1}| = |\overline{a + b + c}| = |a + b + c|$. □

Lemma 4.8. *If $X + X^2 + Y = 0$ 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 $\det(M_{12}) = X^6 - 6X^5 - X^4Y + 16X^4 + 11X^3Y - 10X^3 + 5X^2Y^2 - 7X^2Y - X^2 + XY^3 - 4XY^2 - 3XY - Y^3 - Y^2 = -X^2(X - 1)^5$ and $\text{tr}(M_{12}) = 2X^3-6X^2-3XY-Y^2-Y = X(X-1)^3$ are both nonzero. Since $(1-X)\det(M_{12}) = \text{tr}(M_{12})^2$, we know $\text{Arg}(\det(M_{12})) \neq \text{Arg}(\text{tr}(M_{12})^2)$ and conclude by Lemma 4.2 that the eigenvalues of M_{12} are nonzero and have distinct norm. □

When $X = -1$ no single gadget seems to suffice, but a careful selection of two gadgets together yields a way to handle this case.

Lemma 4.9. *If $X = -1$, $|Y| \neq 2$, and $Y \notin \mathbb{R}$, then the recurrence matrix of unary Gadget 10 has nonzero eigenvalues with distinct norm.*

Proof. Let M_{10} be the recurrence matrix for unary Gadget 10. If M_{10} does not have nonzero eigenvalues with distinct norm, then by Lemma 4.2, $(Y-2)\cdot|-2Y| = \operatorname{tr}(M_{10})\cdot|\det(M_{10})| = \overline{\operatorname{tr}(M_{10})}\cdot\det(M_{10}) = -(\overline{Y}-2)(2Y)$, so $2Y - |Y|^2 - Y\cdot|Y| + 2|Y| = 0$, and $(|Y|-2)(|Y|+Y) = 0$. □

Lemma 4.10. *If $X = -1$, $|Y| = 2$, $Y \notin \mathbb{R}$ and $Y \neq \pm 2i$ then the recurrence matrix of unary Gadget 13 has nonzero eigenvalues with distinct norm.*

Proof. Let M_{13} be the recurrence matrix for unary Gadget 13. Then $\det(M_{13}) = -2Y(Y^2+4) \neq 0$. If the eigenvalues of M_{13} have identical norm then by Lemma 4.2, $\operatorname{tr}(M_{13})^2\overline{\det(M_{13})} = \overline{\operatorname{tr}(M_{13})}^2\det(M_{13})$. Calculation yields $\operatorname{tr}(M_{13})^2\overline{\det(M_{13})} - \overline{\operatorname{tr}(M_{13})}^2\det(M_{13}) = 2(Y-\overline{Y})(\overline{Y}^2+4)^2(\overline{Y}^2+Y^2+8)$, which is nonzero unless $Y \in \mathbb{R}$, $Y = \pm 2i$, or $\overline{Y}^2 + Y^2 + 8 = 0$, and in this last case $\Re(Y^2) = \frac{\overline{Y}^2+Y^2}{2} = -4$ and since $|Y^2| = 4$, $Y^2 = -4$. □

The $4(X-1)^2(X+1) = (Y+2)^2$ condition is somewhat resilient to individual unary recursive gadgets, but by using a second Eigenvalue Shifted Pair, we can reduce it to simpler conditions.

Lemma 4.11. *Suppose $4(X-1)^2(X+1) = (Y+2)^2$. Then either unary Gadget 14 or unary Gadget 15 has nonzero eigenvalues with distinct norm, unless either $X^3+2X^2+X+2Y = 0$, or $X^3+4X^2+2Y-1 = 0$, or both $X, Y \in \mathbb{R}$.*

Proof. The recurrence matrices for unary Gadgets 14 and 15 are

$$M_{14} = \begin{bmatrix} 3a^3+a^6+3ab+b^3 & a+a^4+2a^2b+2b^2+ab^3+b^5 \\ 2a^2+a^5+b+a^3b+2ab^2+b^4 & a^3+3ab+3b^3+b^6 \end{bmatrix}$$

$$M_{15} = \begin{bmatrix} 1+3a^3+a^6+ab+a^2b^2+b^3 & a+a^4+2a^2b+2b^2+ab^3+b^5 \\ 2a^2+a^5+b+a^3b+2ab^2+b^4 & 1+a^3+ab+a^2b^2+3b^3+b^6 \end{bmatrix}$$

so $M_{15} = M_{14} + (X-1)^2 I$. Let α and β be the eigenvalues of M_{14}. Now, $\det(M_{14}) = (X-1)^3(X^3 + 2X^2 + X + 2Y)$ and $\operatorname{tr}(M_{14}) = -2X^3 + 6X + Y^2 + 4Y$ which simplifies to $\operatorname{tr}(M_{14}) = -2X^3 + 6X + Y^2 + 4Y - (Y+2)^2 + 4(X-1)^2(X+1) = -2X(X-1)^2$ using the fact that $4(X-1)^2(X+1) = (Y+2)^2$. Hence $\det(M_{14})$ and $\operatorname{tr}(M_{14})$ are both nonzero (if $X \in \{-1,0,1\}$ then the condition $(Y+2)^2 = 4(X-1)^2(X+1)$ implies $Y \in \mathbb{R}$ as well). Using this same assumption again, $\det(M_{15}) = \det(M_{15}) - (X-1)^2((Y+2)^2 - 4(X-1)^2(X+1)) = (X-1)^3(X^3+4X^2+2Y-1) \neq 0$. Furthermore $\operatorname{tr}[M_{14}]^2 - 4\det(M_{14}) = 4X^2(X-1)^4 - 4(X-1)^3(X^3+2X^2+X+2Y) = -4(X-1)^3(3X^2+X+2Y)$, but if this is zero, then substituting $Y = (-3X^2-X)/2$ into $(Y+2)^2 - 4(X-1)^2(X+1) = 0$ we get $X(X-1)^2(8+9X) = 0$ and $X \in \mathbb{R}$, with $Y \in \mathbb{R}$ as a direct consequence. Corollary 4.2 implies that either Gadget 14 or Gadget 15 has nonzero eigenvalues of distinct norm, unless $\operatorname{tr}(M_{14}) = r(X-1)^2$ and $\det(M_{14}) = s(X-1)^4$ for some $r, s \in \mathbb{R}$. But then $-2X(X-1)^2 = r(X-1)^2$ hence $X = -r/2 \in \mathbb{R}$, and $(X-1)^3(X^3+2X^2+X+2Y) = s(X-1)^4$ hence $Y = (s(X-1) - X - 2X^2 - X^3)/2 \in \mathbb{R}$. □

Now we take advantage of another interesting coincidence; two gadgets with recurrence matrices that have identical trace.

Lemma 4.12. *If $X^2 + X + Y \neq 0$, $4(X-1)^2(X+1) = (Y+2)^2$, and either $X^3 + 2X^2 + X + 2Y = 0$ or $X^3 + 4X^2 + 2Y - 1 = 0$, then the recurrence matrix of unary Gadget 16 or unary Gadget 17 has nonzero eigenvalues with distinct norm, unless both $X, Y \in \mathbb{R}$.*

Proof. The recurrence matrices for unary Gadgets 16 and 17 are

$$M_{16} = \begin{bmatrix} 2a^3 + a^6 + 2ab + a^4b + a^2b^2 + b^3 & a^4 + 3a^2b + b^2 + 2ab^3 + b^5 \\ a^2 + a^5 + 2a^3b + 3ab^2 + b^4 & a^3 + 2ab + a^2b^2 + 2b^3 + ab^4 + b^6 \end{bmatrix}$$

$$M_{17} = \begin{bmatrix} 2a^3 + a^6 + 2ab + a^4b + a^2b^2 + b^3 & a + a^4 + a^2b + b^2 + a^3b^2 + 2ab^3 + b^5 \\ a^2 + a^5 + b + 2a^3b + ab^2 + a^2b^3 + b^4 & a^3 + 2ab + a^2b^2 + 2b^3 + ab^4 + b^6 \end{bmatrix}$$

Let $R = (Y+2)^2 - 4(X-1)^2(X+1)$, $T = X^3 + 2X^2 + X + 2Y$ and $U = X^3 + 4X^2 + 2Y - 1$. Note that regardless of whether $T = 0$ or $U = 0$, $X \in \mathbb{R}$ implies $Y \in \mathbb{R}$, so we will assume $X \notin \mathbb{R}$. The main diagonals of M_{16} and M_{17} are identical, so $\operatorname{tr}(M_{16}) = \operatorname{tr}(M_{17})$. Furthermore, if $T = 0$ then $\operatorname{tr}(M_{16}) = \operatorname{tr}(M_{16}) + R + (X-1)T/2 = X(X-1)^3/2 \neq 0$. If $U = 0$ then $\operatorname{tr}(M_{16}) = \operatorname{tr}(M_{16}) + R + (X-1)U/2 = (X-1)(X^3-1)/2$, and we claim this is nonzero as well. Otherwise, $X^3 = 1$ then since $U = 0$, $Y = -2X^2$ and using $(2+Y)^2 = 4(X-1)^2(X+1)$ we get $(2-2X^2)^2 = 4(X-1)^2(X+1)$ or equivalently $(1-X^2)^2 = (X-1)^2(X+1)$ i.e. $(1-X)^2(1+X)^2 = (X-1)^2(X+1)$ together with $X \notin \mathbb{R}$ we get a contradiction. Next, $\det(M_{17}) = (X-1)^3(X+1)(X^2+X+Y)$ and $\det(M_{16}) = \det(M_{16}) - R(X-1)^2 = (X-1)^3(X+4)(X^2+X+Y)$, so these are both nonzero. If both M_{16} and M_{17} have eigenvalues with equal norm, then applying Lemma 4.2 twice, $\operatorname{Arg}(\det(M_{16})) = \operatorname{Arg}(\operatorname{tr}(M_{16})^2) = \operatorname{Arg}(\operatorname{tr}(M_{17})^2) = \operatorname{Arg}(\det(M_{17}))$. However, this would imply $\operatorname{Arg}(X+4) = \operatorname{Arg}(X+1)$ and $X \in \mathbb{R}$, so we conclude that either M_{16} or M_{17} has nonzero eigenvalues with distinct norm. □

Proof of Lemma 4.7

Proof. By Lemma 4.5, either unary Gadget 9 or 10 has a recurrence matrix with nonzero eigenvalues of distinct norm, except in the following cases.

1. $Y = -2$
2. $X = -1$
3. $X^2 + X + Y = 0$
4. $4(X-1)^2(X+1) = (Y+2)^2$

If $Y = -2$ then $X \notin \mathbb{R}$ and unary Gadget 11 applies by Lemma 4.6. If $X = -1$, then Lemmas 4.9 and 4.10 indicate that either unary Gadget 10 or unary Gadget 13 satisfies the requirement, unless $Y = \pm 2i$. If $X^2 + X + Y = 0$ then $X \notin \mathbb{R}$, lest X and Y are real, so Lemma 4.8 implies that unary Gadget 12 has a recurrence matrix of the required form. Now we may assume $X^2 + X + Y \neq 0$, so by Lemmas 4.11 and 4.12 if $4(X-1)^2(X+1) = (Y+2)^2$ then either unary Gadget 14, 15, 16, or 17 has a recurrence matrix of the required form. □

Proof of Theorem 4.2

Proof. First we transform the problem into coordinates of $X = ab$ and $Y = a^3 + b^3$. We will use binary recursive Gadgets 4, 6, and 7 together with the single-vertex starter gadget given in Figure 1(a) (denote the respective matrices by M_4, M_6, M_7, and S). Calculating the characteristic polynomials

$x^3 + Bx^2 + Cx + D$ of binary recursive Gadgets 4, 6, and 7, we get

$$\begin{aligned}
B_4 &= -(X+Y+1) \\
C_4 &= (X^2+X+Y)(X-1) \\
D_4 &= -X(X-1)^3 \\
B_6 &= -(-2X^3+4X^2+2XY+2X+Y^2+2Y) \\
C_6 &= (X-1)(X^5-4X^4-X^3Y+6X^3+7X^2Y+4X^2+4XY^2+5XY+X+Y^3+2Y^2+Y) \\
D_6 &= -(X-1)^3(2X+Y)(X^4-X^3+X^2Y+3X^2+2XY+X+Y^2+Y) \\
B_7 &= -(-2X^3+2X^2+2X+Y^2+4Y+2) \\
C_7 &= (X-1)^2(X^4-2X^3+2X^2+4XY+6X+2Y^2+4Y+1) \\
D_7 &= -2(X-1)^6X(1+X)
\end{aligned}$$

Condition 2 of Theorem 1 is satisfied with respect to binary recursive Gadget 4 because $\det[S, M_4S, M_4^2S] = (X-1)^4(b^3-a^3) \neq 0$. Now, $\det[S, M_6S, M_6^2S] = (X-1)^5(b^3-a^3)(X^2+X+Y)(X+Y+1)$ and $\det[S, M_7S, M_7^2S] = (X-1)^5(b^3-a^3)(X^2Y+4X^2+2XY+Y^2+Y)$, and since these may be zero for some settings of X and Y, they will need to be encoded into the query. The query is a disjunct of the following conditions on the left hand side.

$$\begin{aligned}
X = 1 &\iff ab = 1 \\
4X^3 = Y^2 &\iff a^3 = b^3 \\
X = -1 \wedge Y = 0 &\iff a^6 = 1 \wedge ab = -1 \\
D_4B_4^3 + C_4^3 \neq 0 &\implies \text{Gadget 4 satisfies Theorem 3.1} \\
(D_6B_6^3 + C_6^3)(X^2+X+Y)(X+Y+1) \neq 0 &\implies \text{Gadget 6 satisfies Theorem 3.1} \\
(D_7B_7^3 + C_7^3)(X^2Y+4X^2+2XY+Y^2+Y) \neq 0 &\implies \text{Gadget 7 satisfies Theorem 3.1}
\end{aligned}$$

Note that we are using Lemma 4.3 to satisfy condition 1 of Theorem 3.1. Using symbolic computation via the CylindricalDecomposition function from Mathematica, we verify that for any $X, Y \in \mathbb{R}$, at least one of the above conditions is true, and we are done. □

Problems of the form $\mathrm{Hol(a,b)}$ where $a^3 = b^3$

Here we characterize all problems of the form $\#[a,1,b] \mid [1,0,0,1]$ where $a^3 = b^3$. It turns out that these are the problems which are #P-hard in general, and yet can be solved in polynomial time when restricted to planar graphs (aside from a few exceptional cases where it is polytime computable in general). We still consider general (not necessarily planar) graphs in this section. The relevant interpolation results can be optained entirely with Gadget 4, using a technique demonstrated in [4].

Lemma 4.13. *The problem* $\#[a,1,a] \mid [1,0,0,1]$ *is* #P*-hard, unless* $a \in \{0, 1, -1, i, -i\}$*, in which case it is in P.*

Proof. If $a \in \mathbb{R}$ then this is already known [5]; a polynomial time algorithm for $a = \pm i$ is in [3]. Now assume $a \notin \mathbb{R}$ and $a \neq \pm i$. Since these problems have an extra degree of symmetry, we use a 2 by 2 recurrence matrix to describe the recursive construction which consists of a single-vertex starter gadget (Figure 1(a)) followed by some number of applications of binary recursive Gadget 4 (no finisher gadget is used here). That is, if $\mathcal{F}$-gate N_i has signature $[a_i, b_i, a_i]$, then the signature of N_{i+1}

is given by $[a_{i+1}, b_{i+1}, a_{i+1}]$ where $[a_{i+1}, b_{i+1}]^{\mathrm{T}} = M[a_i, b_i]^{\mathrm{T}}$, and $M = \begin{bmatrix} a(1+a^2) & 2a \\ 2a^2 & 1+a^2 \end{bmatrix}$. Now, $\det(M) = a(a-1)^2(a+1)^2$ and $\operatorname{tr}(M) = (a+1)(a^2+1)$ are both nonzero under our assumptions. It can be verified (using the Resolve function of Mathematica$^{\mathrm{TM}}$) that $\operatorname{tr}(M)|\det(M)| \neq \overline{\operatorname{tr}(M)}\det(M)$ provided that $a \notin \mathbb{R}$, so by Lemma 4.2, the eigenvalues of M have distinct norm. Also, $M \begin{bmatrix} a \\ 1 \end{bmatrix} = (1+a) \begin{bmatrix} a(2-a+a^2) \\ (1-a+2a^2) \end{bmatrix}$, so the starter gadget is not an eigenvector of M. We conclude by an analogous version of Theorem 3.2 that the problem is #P-hard. □

Lemma 4.14. *If $a^3 = b^3$, then the problem $\#[a,1,b] \mid [1,0,0,1]$ is #P-hard unless $ab \in \{0, 1, -1\}$, in which case it is in P.*

Proof. If $ab = 0$ then $a = b = 0$ and the problem is in P by Theorem 4.1. Otherwise, $ab \neq 0$, let $\omega = ba^{-1}$, and applying a holographic reduction to $\#[a,1,b] \mid [1,0,0,1]$ under the basis $\begin{bmatrix} \omega & 0 \\ 0 & \omega^2 \end{bmatrix}$ we see that the problem $\#[a,1,b] \mid [1,0,0,1]$ is equivalent to $\#[\omega^2 a, 1, \omega b] \mid [1,0,0,1]$, because $\omega^3 = 1$. Since $\omega^2 a = \omega b$, we can apply Lemma 4.13 and the problem $\#[a,1,b] \mid [1,0,0,1]$ is in P if $\pm 1 = \omega^2 a \cdot \omega b = ab$ and #P-hard otherwise. □

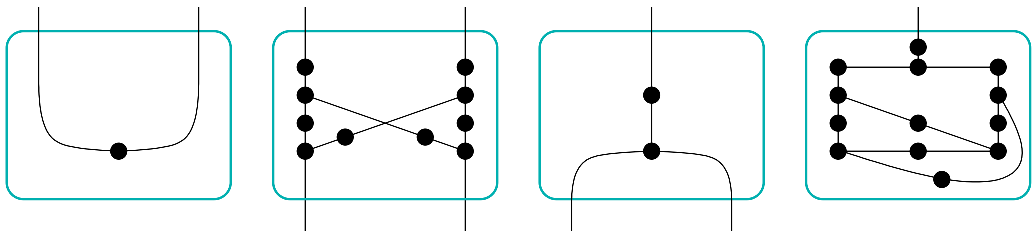

(a) A starter gadget (b) A recursive gadget (c) A finisher gadget (d) A planar embedding of a single iteration

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

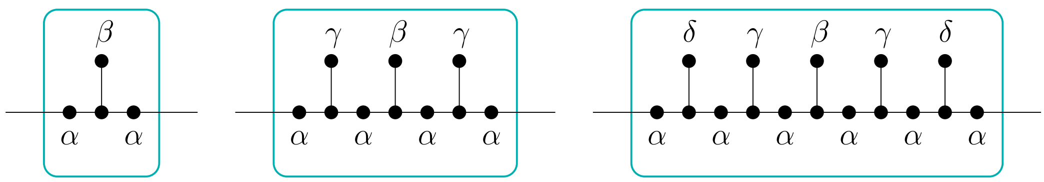

(a) Gadget 1 (b) Gadget 2 (c) Gadget 3

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

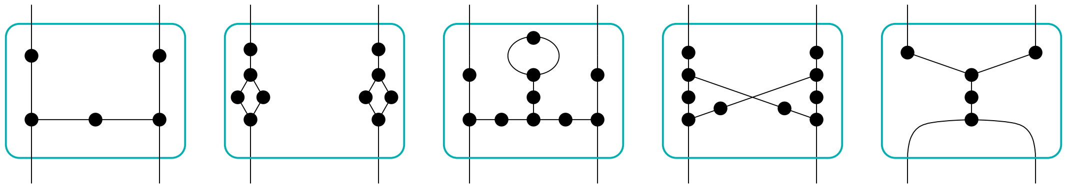

(a) Gadget 4 (b) Gadget 5 (c) Gadget 6 (d) Gadget 7 (e) Gadget 8

Figure 3: Binary recursive gadgets

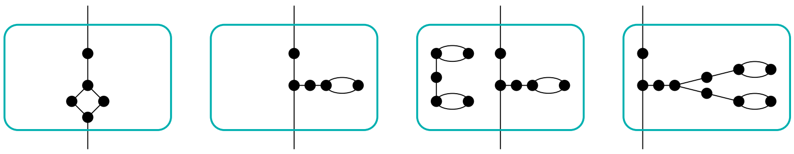

(a) Gadget 9 (b) Gadget 10 (c) Gadget 11 (d) Gadget 12

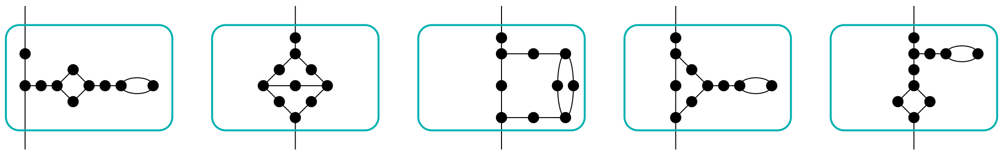

(e) Gadget 13 (f) Gadget 14 (g) Gadget 15 (h) Gadget 16 (i) Gadget 17

Figure 4: Unary recursive gadgets