Holographic Algorithms Beyond Matchgates*

Jin-Yi Cai, Heng Guo, and Tyson Williams

University of Wisconsin–Madison, Madison, WI, USA
{jyc, hguo, tdw}@cs.wisc.edu

Abstract. Holographic algorithms based on matchgates were introduced by Valiant. These algorithms run in polynomial-time and are intrinsically for planar problems. We introduce two new families of holographic algorithms, which work over general, i.e., not necessarily planar, graphs. The two underlying families of constraint functions are of the *affine* and *product* types. These play the role of Kasteleyn's algorithm for counting planar perfect matchings. The new algorithms are obtained by transforming a problem to one of these two families by holographic reductions. We present a polynomial-time algorithm to decide if a given counting problem has a holographic algorithm using these constraint families. When the constraints are symmetric, we give a polynomial-time decision procedure in the size of the succinct presentation of symmetric constraint functions. This procedure shows that the recent dichotomy theorem for Holant problems with symmetric constraints is polynomial-time decidable.

1 Introduction

Recently a number of complexity dichotomy theorems have been obtained for counting problems. Typically, such dichotomy theorems assert that a vast majority of problems expressible within the framework are #P-hard, however an intricate subset manages to escape this fate. They exhibit a great deal of mathematical structure, which leads to a polynomial time algorithm. In recent dichotomy theorems, a pattern has emerged [14,19,21,15,34,23,11,32]. Some of the tractable cases are expressible as "those problems for which there exists a *holographic algorithm*." However, this understanding has been largely restricted to problems where the local constraint functions are symmetric over the Boolean domain. In order to gain a better understanding, we must determine the full extent of holographic algorithms, beyond the symmetric constraints.

Holographic algorithms were first introduced by Valiant [44,43]. They are applicable for any problem that can be expressed as the contraction of a tensor network. Valiant's algorithms have two main ingredients. The first ingredient is to encode computation in planar graphs using matchgates [42,41,9,17,10]. The result of the computation is then obtained by counting the number of perfect matchings in a related planar graph, which can be done in polynomial time by Kasteleyn's (a.k.a. the FKT) algorithm [35,40,36]. The second ingredient is a

* Full version with proofs available at [12].

holographic reduction, which is achieved by a choice of linear basis vectors. The computation can be carried out in any basis since the output of the computation is independent of the basis.

In this paper, we introduce two new families of holographic algorithms. These algorithms holographically reduce to problems expressible by either the *affine* type or the *product* type of constraint functions. Both types of problems are tractable over general (i.e. not necessarily planar) graphs [25], so the holographic algorithms are all polynomial time algorithms and work over general graphs. We present a polynomial time algorithm to decide if a given counting problem has a holographic algorithm over general graphs using the affine or product-type constraint functions. Our algorithm also finds a holographic algorithm when one exists. To formally state this result, we briefly introduce some notation.

The counting problems we consider are those expressible as a Holant problem [24,22,20,25]. A Holant problem is defined by a set $\mathcal{F}$ of constraint functions, which we call signatures, and is denoted by Holant$(\mathcal{F})$. An instance to Holant$(\mathcal{F})$ is a tuple $\Omega = (G, \mathcal{F}, \pi)$ called a signature grid, where $G = (V, E)$ is a graph and π labels each vertex $v \in V$ and its incident edges with some $f_v \in \mathcal{F}$ and its input variables. Here f_v maps $\{0,1\}^{\deg(v)}$ to $\mathbb{C}$. We consider all possible 0-1 edge assignments. An assignment σ to the edges E gives an evaluation $\prod_{v \in V} f_v(\sigma|_{E(v)})$, where $E(v)$ denotes the incident edges of v and $\sigma|_{E(v)}$ denotes the restriction of σ to $E(v)$. The counting problem on the instance Ω is to compute

$$\text{Holant}_\Omega = \sum_{\sigma: E \to \{0,1\}} \prod_{v \in V} f_v\left(\sigma|_{E(v)}\right). \tag{1}$$

For example, consider the problem of counting Perfect Matching on G. This problem corresponds to attaching the Exact-One function at every vertex of G. The Exact-One function is an example of a symmetric signature, which are functions that only depend on the Hamming weight of the input. We denote a symmetric signature by $f = [f_0, f_1, \ldots, f_n]$ where f_w is the value of f on inputs of Hamming weight w. For example, $[0,1,0,0]$ is the Exact-One function on three bits. The output is 1 if and only if the input is 001, 010, or 100, and the output is 0 otherwise.

Holant problems contain both counting constraint satisfaction problems and counting graph homomorphisms as special cases. All three classes of problems have received considerable attention, which has resulted in a number of dichotomy theorems (see [38,33,28,2,27,5,30,8] and [4,3,26,1,25,7,13,29,31,14,6]). Despite this success with #CSP and graph homomorphisms, the case with Holant problems is more difficult. A recent dichotomy theorem for Holant problems with symmetric signatures was obtained in [11]. But the general (i.e. not necessarily symmetric) case has a richer and more intricate structure. The same dichotomy for general signatures remains open. Our first main result makes a solid step forward in understanding holographic algorithms based on affine and product-type signatures in this more difficult setting.

Theorem 1. *There is a polynomial time algorithm to decide, given a finite set of signatures $\mathcal{F}$, whether* $\operatorname{Holant}(\mathcal{F})$ *admits a holographic algorithm based on affine or product-type signatures.*

These holographic algorithms for $\operatorname{Holant}(\mathcal{F})$ are all polynomial time in the size of the problem input Ω. The polynomial time decision algorithm of Theorem 1 is on another level; it decides based on any specific set of signatures $\mathcal{F}$ whether the counting problem $\operatorname{Holant}(\mathcal{F})$ defined by $\mathcal{F}$ has such a holographic algorithm.

However, symmetric signatures are an important special case. Because symmetric signatures can be presented exponentially more succinctly, we would like the decision algorithm to be efficient when measured in terms of this succinct presentation. An algorithm for this case needs to be exponentially faster than the one in Theorem 1. In Theorem 2, we present a polynomial time algorithm for the case of symmetric signatures. The increased efficiency is based on several signature invariants under orthogonal transformations.

Theorem 2. *There is a polynomial time algorithm to decide, given a finite set of symmetric signatures $\mathcal{F}$ expressed in the succinct notation, whether* $\operatorname{Holant}(\mathcal{F})$ *admits a holographic algorithm based on affine or product-type signatures.*

A dichotomy theorem classifies every set of signatures as defining either a tractable problem or an intractable problem (e.g. #P-hard). Yet it would be more useful if given a specific set of signatures, one could decide to which case it belongs. This is the decidability problem of a dichotomy theorem. In [11], a dichotomy regarding symmetric complex-weighted signatures for Holant problem was proved. However, the decidability problem was left open. Of the five tractable cases in this dichotomy theorem, three of them are easy to decide, but the remaining two cases are more challenging, which are (1) holographic algorithms using affine signatures and (2) holographic algorithms using product-type signatures. As a consequence of Theorem 2, this decidability is now proved.

Corollary 3. *The dichotomy theorem for symmetric complex-weighted Holant problems in [11] is decidable in polynomial time.*

Previous work on holographic algorithms focused almost exclusively on those with matchgates [44,43,16,19,17,18,32]. (This has led to a misconception in the community that holographic algorithms are always based on matchgates.) The first example of a holographic algorithm using something other than matchgates came in [24]. These holographic algorithms use generalized Fibonacci gates. A symmetric signature $f = [f_0, f_1, \ldots, f_n]$ is a generalized Fibonacci gate of type $\lambda \in \mathbb{C}$ if $f_{k+2} = \lambda f_{k+1} + f_k$ holds for all $k \in \{0, 1, \ldots, n-2\}$. The standard Fibonacci gates are of type $\lambda = 1$, in which case, the entries of the signature satisfy the recurrence relation of the Fibonacci numbers. The generalized Fibonacci gates were immediately put to use in a dichotomy theorem [22]. As it turned out, for nearly all values of λ, the generalized Fibonacci gates are holographically equivalent to product-type signatures. However, generalized Fibonacci gates are symmetric by definition. A main contribution of this paper is to extend the reach of holographic algorithms, other than those based on matchgates, beyond the symmetric case.

The constraint functions we call signatures are essentially tensors. Our central object of study can be rephrased as the orbits of affine and product-type tensors when acted upon by the orthogonal group (and related groups). We show that one can efficiently decide if any such orbit of a given tensor intersects the set of affine or product-type tensors. This result also generalizes to a set of tensors as stated in Theorems 1 and 2. In contrast, this orbit problem with the general linear group acting on two arbitrary tensors is NP-hard [37]. The so-called orbit closure problem has a fundamental importance in the foundation of geometric complexity theory [39].

Our techniques are mainly algebraic. A particularly important insight is that an orthogonal transformation in the standard basis is equivalent to a diagonal transformation in the $\left[\begin{smallmatrix} 1 & 1 \\ i & -i \end{smallmatrix}\right]$ basis, a type of correspondence as in Fourier transform. Since diagonal transformations are much easier to understand, this gives us a great advantage in understanding orbits under orthogonal transformations. Also, the groups of transformations that stabilize the affine and product-type signatures play an important role in our proofs.

2 Preliminaries

The framework of Holant problems is defined for functions mapping any $[q]^k \to \mathbb{F}$ for a finite q and some field $\mathbb{F}$. In this paper, we investigate some of the tractable complex-weighted Boolean Holant problems, that is, all functions are $[2]^k \to \mathbb{C}$. Strictly speaking, for consideration of models of computation, functions take complex algebraic numbers.

A *signature grid* $\Omega = (G, \mathcal{F}, \pi)$ consists of a graph $G = (V, E)$, where each vertex is labeled by a function $f_v \in \mathcal{F}$, and $\pi : V \to \mathcal{F}$ is the labeling. The Holant problem on instance Ω is to evaluate $\text{Holant}_\Omega = \sum_\sigma \prod_{v \in V} f_v(\sigma\mid_{E(v)})$, a sum over all edge assignments $\sigma : E \to \{0, 1\}$. A function f_v can be represented by listing its values in lexicographical order as in a truth table, which is a vector in $\mathbb{C}^{2^{\deg(v)}}$, or as a tensor in $(\mathbb{C}^2)^{\otimes \deg(v)}$. We also use $f_{\mathbf{x}}$ to denote the value $f(\mathbf{x})$, where $\mathbf{x}$ is a binary string. A function $f \in \mathcal{F}$ is also called a *signature*. A symmetric signature f on k Boolean variables can be expressed as $[f_0, f_1, \ldots, f_k]$, where f_w is the value of f on inputs of Hamming weight w. A signature f of arity n is *degenerate* if there exist unary signatures $u_j \in \mathbb{C}^2$ $(1 \le j \le n)$ such that $f = u_1 \otimes \cdots \otimes u_n$. A symmetric degenerate signature has the form $u^{\otimes n}$.

A Holant problem is parametrized by a set of signatures.

Definition 4. *Given a set of signatures $\mathcal{F}$, we define* Holant$(\mathcal{F})$ *as:*
Input: A signature grid $\Omega = (G, \mathcal{F}, \pi)$*;*
Output: Holant_Ω.

To introduce the idea of holographic reductions, it is convenient to consider bipartite graphs. For a general graph, we can always transform it into a bipartite graph while preserving the Holant value, as follows. For each edge in the graph, we replace it by a path of length two. (This operation is called the *2-stretch* of the graph and yields the edge-vertex incidence graph.) Each new vertex is assigned the binary EQUALITY signature $(=_2) = [1, 0, 1]$. We use Holant $(\mathcal{F} \mid \mathcal{G})$ to denote the Holant problem on bipartite graphs $H = (U, V, E)$, where each vertex in U or V is assigned a signature in $\mathcal{F}$ or $\mathcal{G}$, respectively. An instance for this bipartite Holant problem is a bipartite signature grid denoted by $\Omega = (H;\ \mathcal{F} \mid \mathcal{G};\ \pi)$. Signatures in $\mathcal{F}$ are considered as row vectors (or covariant tensors); signatures in $\mathcal{G}$ are considered as column vectors (or contravariant tensors).

For a 2-by-2 matrix T and a signature set $\mathcal{F}$, define $T\mathcal{F} = \{g \mid \exists f \in \mathcal{F}$ of arity $n,\ g = T^{\otimes n} f\}$, similarly for $\mathcal{F}T$. Whenever we write $T^{\otimes n} f$ or $T\mathcal{F}$, we view the signatures as column vectors; similarly for $fT^{\otimes n}$ or $\mathcal{F}T$ as row vectors. Let T be an element of $\mathbf{GL}_2(\mathbb{C})$, the group of invertible 2-by-2 complex matrices. The holographic transformation defined by T is the following operation: given a signature grid $\Omega = (H;\ \mathcal{F} \mid \mathcal{G};\ \pi)$, for the same graph H, we get a new grid $\Omega' = (H;\ \mathcal{F}T \mid T^{-1}\mathcal{G};\ \pi')$ by replacing each signature in $\mathcal{F}$ or $\mathcal{G}$ with the corresponding signature in $\mathcal{F}T$ or $T^{-1}\mathcal{G}$.

Theorem 5 (Valiant's Holant Theorem [44]). *If there is a holographic transformation mapping signature grid Ω to Ω', then* $\mathrm{Holant}_{\Omega} = \mathrm{Holant}_{\Omega'}$.

Therefore, an invertible holographic transformation does not change the complexity of the Holant problem in the bipartite setting. Furthermore, there is a particular kind of holographic transformation, the orthogonal transformation, that preserves the binary equality and thus can be used freely in the standard setting. Let $\mathbf{O}_2(\mathbb{C})$ be the group of 2-by-2 complex matrices that are orthogonal. Recall that a matrix T is orthogonal if $TT^{\mathsf{T}} = I$. We also use $\mathbf{SO}_2(\mathbb{C})$ to denote the group of special orthogonal matrices, i.e. the subgroup of $\mathbf{O}_2(\mathbb{C})$ with determinant 1.

The following two signature sets are tractable without a holographic transformation [25].

Definition 6. *A k-ary function $f(x_1, \ldots, x_k)$ is* affine *if it has the form $\lambda \cdot \chi_{Ax=0} \cdot i^{\sum_{j=1}^{n} \langle \mathbf{v}_j, x \rangle}$, where $\lambda \in \mathbb{C}$, $x = (x_1, x_2, \ldots, x_k, 1)^T$, A is a matrix over $\mathbb{F}_2$, $\mathbf{v}_j$ is a vector over $\mathbb{F}_2$, and χ is a 0-1 indicator function such that $\chi_{Ax=0}$ is 1 iff $Ax = 0$. Note that the dot product $\langle \mathbf{v}_j, x \rangle$ is calculated over $\mathbb{F}_2$, while the summation $\sum_{j=1}^{n}$ on the exponent of $i = \sqrt{-1}$ is evaluated as a sum mod 4 of 0-1 terms. We use $\mathscr{A}$ to denote the set of all affine functions.*

An equivalent way to express the exponent of i is as a quadratic polynomial where all cross terms have an even coefficient.

Definition 7. *A function is of* product type *if it can be expressed as a product of unary functions, binary equality functions* $([1, 0, 1])$*, and binary disequality functions* $([0, 1, 0])$*. We use $\mathscr{P}$ to denote the set of product-type functions.*

The tractable sets $\mathscr{A}$ and $\mathscr{P}$ are still tractable under a suitable holographic transformation. This is captured by the following definition.

Definition 8. *A set $\mathcal{F}$ of signatures is $\mathscr{A}$-transformable (resp. $\mathscr{P}$-transformable) if there exists a holographic transformation T such that $\mathcal{F} \subseteq T\mathscr{A}$ (resp. $\mathcal{F} \subseteq T\mathscr{P}$) and $[1,0,1]T^{\otimes 2} \in \mathscr{A}$ (resp. $[1,0,1]T^{\otimes 2} \in \mathscr{P}$).*

To refine the above definition, we consider the stabilizer group of $\mathscr{A}$, which is $\operatorname{Stab}(\mathscr{A}) = \{T \in \mathbf{GL}_2(\mathbb{C}) \mid T\mathscr{A} \subseteq \mathscr{A}\}$. Technically this set is the left stabilizer group of $\mathscr{A}$, but it turns out that the left and right stabilizer groups of $\mathscr{A}$ coincide. Let $D = \left[\begin{smallmatrix} 1 & 0 \\ 0 & i \end{smallmatrix}\right]$ and $H_2 = \frac{1}{\sqrt{2}}\left[\begin{smallmatrix} 1 & 1 \\ 1 & -1 \end{smallmatrix}\right]$. Also let $X = \left[\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right]$ and $Z = \frac{1}{\sqrt{2}}\left[\begin{smallmatrix} 1 & 1 \\ i & -i \end{smallmatrix}\right]$. Note that $Z = DH_2$ and that $D^2 Z = \frac{1}{\sqrt{2}}\left[\begin{smallmatrix} 1 & 1 \\ -i & i \end{smallmatrix}\right] = ZX$, hence $X = Z^{-1}D^2 Z$. It is easy to verify that $D, H_2, X, Z \in \operatorname{Stab}(\mathscr{A})$. In fact, $\operatorname{Stab}(\mathscr{A}) = \mathbb{C}^* \cdot \langle D, H_2 \rangle$, i.e. all nonzero scalar multiples of the group generated by D and H_2. Throughout the paper, we use α to denote $\frac{1+i}{\sqrt{2}} = \sqrt{i} = e^{\frac{\pi i}{4}}$.

Definition 9. *A symmetric signature f of arity n is in, respectively, $\mathscr{A}_1$, or $\mathscr{A}_2$, or $\mathscr{A}_3$ if there exist an $H \in \mathbf{O}_2(\mathbb{C})$ and $c \in \mathbb{C} - \{0\}$ such that f has the form, respectively, $cH^{\otimes n}(\left[\begin{smallmatrix} 1 \\ 1 \end{smallmatrix}\right]^{\otimes n} + \beta \left[\begin{smallmatrix} 1 \\ -1 \end{smallmatrix}\right]^{\otimes n})$, $cH^{\otimes n}(\left[\begin{smallmatrix} 1 \\ i \end{smallmatrix}\right]^{\otimes n} + \left[\begin{smallmatrix} 1 \\ -i \end{smallmatrix}\right]^{\otimes n})$, or $cH^{\otimes n}(\left[\begin{smallmatrix} 1 \\ \alpha \end{smallmatrix}\right]^{\otimes n} + i^r \left[\begin{smallmatrix} 1 \\ -\alpha \end{smallmatrix}\right]^{\otimes n})$ with $\beta = \alpha^{tn+2r}$, $r \in \{0,1,2,3\}$, and $t \in \{0,1\}$.*

The three sets $\mathscr{A}_1$, $\mathscr{A}_2$, and $\mathscr{A}_3$ capture all symmetric $\mathscr{A}$-transformable signatures. For $i \in \{1,2,3\}$, when such an orthogonal H exists, we say that $f \in \mathscr{A}_i$ with transformation H.

Lemma 10 (Lemma 8.10 in full version of [11]). *Let f be a non-degenerate symmetric signature. Then f is $\mathscr{A}$-transformable iff $f \in \mathscr{A}_1 \cup \mathscr{A}_2 \cup \mathscr{A}_3$.*

We have a similar characterization for $\mathscr{P}$-transformable signatures using the stabilizer group of $\mathscr{P}$, $\operatorname{Stab}(\mathscr{P}) = \{T \in \mathbf{GL}_2(\mathbb{C}) \mid T\mathscr{P} \subseteq \mathscr{P}\}$. The group $\operatorname{Stab}(\mathscr{P})$ is generated by matrices of the form $\left[\begin{smallmatrix} 1 & 0 \\ 0 & \nu \end{smallmatrix}\right]$ for any $\nu \in \mathbb{C}$ and $X = \left[\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right]$.

Definition 11. *A symmetric signature f of arity n is in $\mathscr{P}_1$ if there exist an $H \in \mathbf{O}_2(\mathbb{C})$ and a nonzero $c \in \mathbb{C}$ such that $f = cH^{\otimes n}\left(\left[\begin{smallmatrix} 1 \\ 1 \end{smallmatrix}\right]^{\otimes n} + \beta \left[\begin{smallmatrix} 1 \\ -1 \end{smallmatrix}\right]^{\otimes n}\right)$, where $\beta \neq 0$.*

It is easy to check that $\mathscr{A}_1 \subset \mathscr{P}_1$. We define $\mathscr{P}_2 = \mathscr{A}_2$. Similarly, for $i \in \{1,2\}$, when such an H exists, we say that $f \in \mathscr{P}_i$ with transformation H. The following lemma is similar to Lemma 10.

Lemma 12 (Lemma 8.13 in full version of [11]). *Let f be a non-degenerate symmetric signature. Then f is $\mathscr{P}$-transformable iff $f \in \mathscr{P}_1 \cup \mathscr{P}_2$.*

3 General Signatures

In this section we consider general (i.e. not necessarily symmetric) signatures. Let f be a signature of arity n. It is given as a column vector in $\mathbb{C}^{2^n}$ with bit length $N = O(2^n)$. We denote its entries by $f_{\mathbf{x}} = f(\mathbf{x})$ indexed by $\mathbf{x} \in \{0,1\}^n$.

The entries are from a fixed degree algebraic extension of $\mathbb{Q}$ and we may assume arithmetic operations take unit time.

We begin with $\mathscr{A}$-transformable signatures. Let f be a signature and $H = \left[\begin{smallmatrix} a & b \\ -b & a \end{smallmatrix}\right] \in \mathbf{SO}_2(\mathbb{C})$ where $a^2 + b^2 = 1$. Notice that $v_0 = (1, i)$ and $v_1 = (1, -i)$ are row eigenvectors of H with eigenvalues $a - bi$ and $a + bi$ respectively.

For a vector $\mathbf{u} = (u_1, \ldots, u_n) \in \{0,1\}^n$ of length n, let $v_{\mathbf{u}} = v_{u_1} \otimes v_{u_2} \otimes \ldots \otimes v_{u_n}$, and let $w(\mathbf{u})$ be the Hamming weight of $\mathbf{u}$. Then for the 2^n-by-2^n matrix $H^{\otimes n}$, $v_{\mathbf{u}}$ is a row eigenvector with eigenvalue $(a-bi)^{n-w(\mathbf{u})}(a+bi)^{w(\mathbf{u})} = (a-bi)^{n-2w(\mathbf{u})} = (a+bi)^{2w(\mathbf{u})-n}$ as $(a+bi)(a-bi) = a^2+b^2 = 1$. Let $Z' = \left[\begin{smallmatrix} 1 & i \\ 1 & -i \end{smallmatrix}\right]$ and $\hat{f} = Z'^{\otimes n} f$. Then $\hat{f}_{\mathbf{u}} = \langle v_{\mathbf{u}}, f \rangle$, as a dot product. The following lemma summarizes the above discussion and is a very important ingredient of this paper. It states that proper orthogonal transformations are diagonal transformations in the $\left[\begin{smallmatrix} 1 & i \\ 1 & -i \end{smallmatrix}\right]$ basis.

Lemma 13. *Suppose f and g are signatures of arity n and let $H = \left[\begin{smallmatrix} a & b \\ -b & a \end{smallmatrix}\right]$ and $T = \left[\begin{smallmatrix} a-bi & 0 \\ 0 & a+bi \end{smallmatrix}\right]$. Then $g = H^{\otimes n} f$ iff $\hat{g} = T^{\otimes n} \hat{f}$.*

With Lemma 13, we characterize signatures that are invariant under $\mathbf{SO}_2(\mathbb{C})$ transformations.

Lemma 14. *Let f be a signature. Then f is invariant under transformations in $\mathbf{SO}_2(\mathbb{C})$ (up to a nonzero constant) iff the support of $\hat{f}$ contains at most one Hamming weight.*

With Lemma 13 and Lemma 14, we are able to give the algorithm for $\mathscr{A}$-transformable signatures.

Theorem 15. *There is a polynomial time algorithm to decide, for any finite set of signatures $\mathcal{F}$, whether $\mathcal{F}$ is $\mathscr{A}$-transformable. If so, at least one transformation can be found.*

The algorithm for $\mathscr{P}$ is also based on Lemma 13. The difference here is that we need to first factor the signatures. We show a unique factorization lemma for signatures in general.

Definition 16. *We call a function f of arity n on variable set $\mathbf{x}$* reducible *if there exist f_1 and f_2 of arities n_1 and n_2 on variable sets $\mathbf{x}_1$ and $\mathbf{x}_2$, respectively, such that $1 \le n_1, n_2 \le n-1$, $\mathbf{x}_1 \cup \mathbf{x}_2 = \mathbf{x}$, $\mathbf{x}_1 \cap \mathbf{x}_2 = \emptyset$, and $f(\mathbf{x}) = f_1(\mathbf{x}_1) f_2(\mathbf{x}_2)$. Otherwise we call f* irreducible.

If a function f is reducible, then we can factor it into functions of smaller arity. This procedure can be applied recursively and terminates when all components are irreducible. Therefore any function has at least one irreducible factorization. We show that such a factorization is unique for functions that are not identically zero. Furthermore, it can be computed in polynomial time.

Lemma 17. *Let f be a function of arity n on variables $\mathbf{x}$ that is not identically zero. Assume there exist irreducible functions f_i and g_j, and two partitions $\{\mathbf{x}_i\}$ and $\{\mathbf{y}_j\}$ of $\mathbf{x}$ for $1 \le i \le k$ and $1 \le j \le k'$, such that $f(\mathbf{x}) = \prod_{i=1}^{k} f_i(\mathbf{x}_i) = \prod_{j=1}^{k'} g_j(\mathbf{y}_j)$. Then $k = k'$, the partitions are the same, and $\{f_i\}$ and $\{g_j\}$ are the same up to a permutation.*

The factorization algorithm leads to a decision algorithm for membership in $\mathscr{P}$. Combined with Lemma 13, we can obtain the algorithm for $\mathscr{P}$-transformable signatures.

Theorem 18. *There is a polynomial time algorithm to decide, for any finite set of signatures $\mathcal{F}$, whether $\mathcal{F}$ is $\mathscr{P}$-transformable. If so, at least one transformation can be found.*

4 Symmetric Signatures

In this section, we consider the case when the signatures are symmetric. The significant difference is that a symmetric signature of arity n is given by $n+1$ values, instead of 2^n values. This exponentially more succinct representation requires us to find a more efficient algorithm. To begin, we provide efficient algorithms to decide membership each of $\mathscr{A}_1$, $\mathscr{A}_2$, and $\mathscr{A}_3$ for a single signature. If the signature is in one of the sets, then the algorithm also finds at least one corresponding orthogonal transformation. By Lemma 10, this is enough to check if a signature is $\mathscr{A}$-transformable.

We say a signature f satisfies a second order recurrence relation, if for all $0 \le k \le n-2$, there exist $a, b, c \in \mathbb{C}$ not all zero, such that $af_k + bf_{k+1} + cf_{k+2} = 0$. In fact, satisfying a second order recurrence relation with $b^2 - 4ac \neq 0$ is a necessary condition for a signature to be $\mathscr{A}$- or $\mathscr{P}$-transformable. This also implies a tensor decomposition of f. The following definition of the θ function is crucial.

Definition 19. *For a pair of linearly independent vectors $v_0 = \left[\begin{smallmatrix} a_0 \\ b_0 \end{smallmatrix}\right]$ and $v_1 = \left[\begin{smallmatrix} a_1 \\ b_1 \end{smallmatrix}\right]$, we define $\theta(v_0, v_1) = \left(\frac{a_0 a_1 + b_0 b_1}{a_1 b_0 - a_0 b_1}\right)^2$. Furthermore, suppose that a signature f of arity $n \ge 3$ can be expressed as $f = v_0^{\otimes n} + v_1^{\otimes n}$, where v_0 and v_1 are linearly independent. Then we define $\theta(f) = \theta(v_0, v_1)$.*

Intuitively, this formula is the square of the cotangent of the angle from v_0 to v_1. This notion of cotangent is properly extended to the complex domain. By insisting that v_0 and v_1 be linearly independent, we ensure $\theta(v_0, v_1)$ is well-defined. The expression is squared so that $\theta(v_0, v_1) = \theta(v_1, v_0)$. Let $f = v_0^{\otimes n} + v_1^{\otimes n}$ be a non-degenerate signature of arity $n \ge 3$. Since f is non-degenerate, v_0 and v_1 are linearly independent. This expression for f via v_0 and v_1 is unique to up a root of unity. In particular, $\theta(f)$ from Definition 19 is well-defined since every possible expression gives the same value for θ. It is easy to verify that θ is invariant under an orthogonal transformation. Formally, we have the following lemma, which is proved by simple algebra.

Lemma 20. *For two linearly independent vectors v_0, $v_1 \in \mathbb{C}^2$ and $H \in \mathbf{O}_2(\mathbb{C})$, let $\widehat{v_0} = Hv_0$ and $\widehat{v_1} = Hv_1$. Then $\theta(v_0, v_1) = \theta(\widehat{v_0}, \widehat{v_1})$.*

Now we have some necessary conditions for a signature f to be in $\mathscr{A}_1 \cup \mathscr{A}_2 \cup \mathscr{A}_3$. First f must satisfy a second order recurrence relation with $b^2 - 4ac \neq 0$. Then $\theta(f)$ is well defined. It is easy to observe $\theta(f) = 0, -1, -\frac{1}{2}$ for f in $\mathscr{P}_1$, $\mathscr{A}_2$, $\mathscr{A}_3$ respectively. Recall that $\mathscr{A}_1 \subseteq \mathscr{P}_1$ and $\mathscr{A}_2 = \mathscr{P}_2$.

This condition via $\theta(f)$ is not sufficient for f to be $\mathscr{A}$-transformable. For example, if $f = v_0^{\otimes n} + v_1^{\otimes n}$ with $v_0 = [1, i]$ and v_1 is not a multiple of $[1, -i]$, then $\theta(f) = -1$ but f is not in $\mathscr{A}_2 = \mathscr{P}_2$. Nevertheless, this is essentially the only exceptional case. The other cases are handled with some extra conditions on the coefficients, as follows.

Lemma 21. *Let $f = v_0^{\otimes n} + v_1^{\otimes n}$ be a symmetric signature of arity $n \geq 3$, where $v_0 = \left[\begin{smallmatrix} a_0 \\ b_0 \end{smallmatrix}\right]$ and $v_1 = \left[\begin{smallmatrix} a_1 \\ b_1 \end{smallmatrix}\right]$ are linearly independent. Then $f \in \mathscr{A}_1$ iff $\theta(f) = 0$ and there exist an $r \in \{0, 1, 2, 3\}$ and $t \in \{0, 1\}$ such that $a_1^n = \alpha^{tn+2r} b_0^n \neq 0$ or $b_1^n = \alpha^{tn+2r} a_0^n \neq 0$.*

Lemma 22. *Let $f = v_0^{\otimes n} + v_1^{\otimes n}$ be a symmetric signature of arity $n \geq 3$, where $v_0 = \left[\begin{smallmatrix} a_0 \\ b_0 \end{smallmatrix}\right]$ and $v_1 = \left[\begin{smallmatrix} a_1 \\ b_1 \end{smallmatrix}\right]$ are linearly independent. Then $f \in \mathscr{A}_3$ iff there exist an $\varepsilon \in \{1, -1\}$ and $r \in \{0, 1, 2, 3\}$ such that $a_1 \left(\sqrt{2} a_0 + \varepsilon i b_0\right) = b_1 \left(\varepsilon i a_0 - \sqrt{2} b_0\right)$, $a_1^n = i^r \left(\varepsilon i a_0 - \sqrt{2} b_0\right)^n$, and $b_1^n = i^r \left(\sqrt{2} a_0 + \varepsilon i b_0\right)^n$.*

For $\mathscr{A}_2 = \mathscr{P}_2$, we require a stronger condition.

Lemma 23 (Lemma 8.8 in full version of [11]). *Let f be a non-degenerate symmetric signature. Then $f \in \mathscr{A}_2$ iff f is of the form $c\left(\left[\begin{smallmatrix} 1 \\ i \end{smallmatrix}\right]^{\otimes n} + \beta \left[\begin{smallmatrix} 1 \\ -i \end{smallmatrix}\right]^{\otimes n}\right)$ for some $c, \beta \neq 0$.*

To summarize, we have the following lemma.

Lemma 24. *Given a non-degenerate symmetric signature f of arity at least 3, there is a polynomial time algorithm to decide whether $f \in \mathscr{A}_k$ for each $k \in \{1, 2, 3\}$. If so, k is unique and at least one corresponding orthogonal transformation can be found in polynomial time.*

Next we show that if a non-degenerate signature f of arity $n \geq 3$ is in $\mathscr{A}_1$, $\mathscr{A}_2$, or $\mathscr{A}_3$, then for any set $\mathcal{F}$ containing f, there are only $O(n)$ many transformations to be checked to decide whether $\mathcal{F}$ is $\mathscr{A}$-transformable.

Lemma 25. *Let $\mathcal{F}$ be a set of symmetric signatures and suppose $\mathcal{F}$ contains a non-degenerate signature $f \in \mathscr{A}_1$ of arity $n \geq 3$ with $H \in \mathbf{O}_2(\mathbb{C})$. Then $\mathcal{F}$ is $\mathscr{A}$-transformable iff $\mathcal{F}$ is a subset of $H\mathscr{A}$, or $H \left[\begin{smallmatrix} 1 & 1 \\ 1 & -1 \end{smallmatrix}\right] \mathscr{A}$, or $H \left[\begin{smallmatrix} 1 & 1 \\ 1 & -1 \end{smallmatrix}\right] \left[\begin{smallmatrix} 1 & 0 \\ 0 & \alpha \end{smallmatrix}\right] \mathscr{A}$.*

Lemma 26. *Let $\mathcal{F}$ be a set of symmetric signatures and suppose $\mathcal{F}$ contains a non-degenerate signature $f \in \mathscr{A}_2$ of arity $n \geq 3$. Then there exists a set $\mathcal{H} \subseteq \mathbf{O}_2(\mathbb{C})$ of size $O(n)$ such that $\mathcal{F}$ is $\mathscr{A}$-transformable iff there exists an $H \in \mathcal{H}$ such that $\mathcal{F} \subseteq H\mathscr{A}$. Moreover $\mathcal{H}$ can be computed in polynomial time in the input length of the symmetric signature f.*

Lemma 27. *Let $\mathcal{F}$ be a set of symmetric signatures and suppose $\mathcal{F}$ contains a non-degenerate signature $f \in \mathscr{A}_3$ of arity $n \geq 3$ with $H \in \mathbf{O}_2(\mathbb{C})$. Then $\mathcal{F}$ is $\mathscr{A}$-transformable iff $\mathcal{F} \subseteq H \left[\begin{smallmatrix} 1 & 0 \\ 0 & \alpha \end{smallmatrix}\right] \mathscr{A}$.*

Now we can decide if a finite set of signatures is $\mathscr{A}$-transformable. To avoid trivialities, we assume $\mathcal{F}$ contains a non-degenerate signature of arity at least 3.

Theorem 28. *There is a polynomial time algorithm to decide, for any finite input set $\mathcal{F}$ of symmetric signatures containing a non-degenerate signature f of arity $n \geq 3$, whether $\mathcal{F}$ is $\mathscr{A}$-transformable.*

Now we consider $\mathscr{P}$-transformable signatures. To decide if a single signature is $\mathscr{P}$-transformable, it is equivalent to decide membership in $\mathscr{P}_1 \cup \mathscr{P}_2$ by Lemma 12. The following lemma tells how to decide the membership of $\mathscr{P}_1$.

Lemma 29. *Let $f = v_0^{\otimes n} + v_1^{\otimes n}$ be a symmetric signature of arity $n \geq 3$, where v_0 and v_1 are linearly independent. Then $f \in \mathscr{P}_1$ iff $\theta(f) = 0$.*

Since $\mathscr{A}_2 = \mathscr{P}_2$, deciding membership in $\mathscr{P}_2$ is handled by Lemma 23. Using Lemma 29 and Lemma 23, we can efficiently decide membership in $\mathscr{P}_1 \cup \mathscr{P}_2$.

Lemma 30. *Given a non-degenerate symmetric signature f of arity at least 3, there is a polynomial time algorithm to decide whether $f \in \mathscr{P}_k$ for some $k \in \{1, 2\}$. If so, k is unique and at least one corresponding orthogonal transformation can be found in polynomial time.*

With a signature in $\mathscr{P}_1 \cup \mathscr{P}_2$, we can decide if a set of symmetric signatures is $\mathscr{P}$-transformable.

Lemma 31. *Let $\mathcal{F}$ be a set of symmetric signatures and suppose $\mathcal{F}$ contains a non-degenerate signature $f \in \mathscr{P}_1$ of arity $n \geq 3$ with $H \in \mathbf{O}_2(\mathbb{C})$. Then $\mathcal{F}$ is $\mathscr{P}$-transformable iff $\mathcal{F} \subseteq H \left[\begin{smallmatrix} 1 & 1 \\ 1 & -1 \end{smallmatrix}\right] \mathscr{P}$.*

Lemma 32. *Let $\mathcal{F}$ be a set of symmetric signatures and suppose $\mathcal{F}$ contains a non-degenerate signature $f \in \mathscr{P}_2$ of arity $n \geq 3$. Then $\mathcal{F}$ is $\mathscr{P}$-transformable iff all non-degenerate signatures in $\mathcal{F}$ are contained in $\mathscr{P}_2 \cup \{=_2\}$.*

With all these results, we show how to decide if a finite set of signatures is $\mathscr{P}$-transformable. To avoid trivialities, we assume $\mathcal{F}$ contains a non-degenerate signature of arity at least 3.

Theorem 33. *There is a polynomial time algorithm to decide, for any finite input set $\mathcal{F}$ of symmetric signatures containing a non-degenerate signature f of arity $n \geq 3$, whether $\mathcal{F}$ is $\mathscr{P}$-transformable.*

Acknowledgements All authors were supported by NSF CCF-1217549. The second author was also supported by a Simons Award for Graduate Students in Theoretical Computer Science from the Simons Foundation.

References

1. Andrei Bulatov, Martin Dyer, Leslie Ann Goldberg, Markus Jalsenius, and David Richerby. The complexity of weighted Boolean #CSP with mixed signs. *Theor. Comput. Sci.*, 410(38-40):3949–3961, 2009.
2. Andrei Bulatov and Martin Grohe. The complexity of partition functions. *Theor. Comput. Sci.*, 348(2):148–186, 2005.
3. Andrei A. Bulatov. The complexity of the counting constraint satisfaction problem. *J. ACM*, 60(5):34:1–34:41, 2013.
4. Andrei A. Bulatov and Víctor Dalmau. Towards a dichotomy theorem for the counting constraint satisfaction problem. *Inform. and Comput.*, 205(5):651–678, 2007.
5. Jin-Yi Cai and Xi Chen. A decidable dichotomy theorem on directed graph homomorphisms with non-negative weights. In *FOCS*, pages 437–446. IEEE Computer Society, 2010.
6. Jin-Yi Cai and Xi Chen. Complexity of counting CSP with complex weights. In *STOC*, pages 909–920. ACM, 2012. arXiv:1111.2384.
7. Jin-Yi Cai, Xi Chen, and Pinyan Lu. Non-negatively weighted #CSP: An effective complexity dichotomy. In *IEEE Conference on Computational Complexity*, pages 45–54. IEEE Computer Society, 2011.
8. Jin-Yi Cai, Xi Chen, and Pinyan Lu. Graph homomorphisms with complex values: A dichotomy theorem. *SIAM J. Comput.*, 42(3):924–1029, 2013.
9. Jin-Yi Cai, Vinay Choudhary, and Pinyan Lu. On the theory of matchgate computations. *Theory of Computing Systems*, 45(1):108–132, 2009.
10. Jin-Yi Cai and Aaron Gorenstein. Matchgates revisited. *CoRR*, abs/1303.6729, 2013. To appear in "Theory of Computing" (ToC) `theoryofcomputing.org`.
11. Jin-Yi Cai, Heng Guo, and Tyson Williams. A complete dichotomy rises from the capture of vanishing signatures (extended abstract). In *STOC*, pages 635–644. ACM, 2013. arXiv:1204.6445.
12. Jin-Yi Cai, Heng Guo, and Tyson Williams. Holographic algorithms beyond matchgates. *CoRR*, abs/1307.7430, 2013. To appear at ICALP 2014.
13. Jin-Yi Cai, Sangxia Huang, and Pinyan Lu. From Holant to #CSP and back: Dichotomy for Holant^c problems. *Algorithmica*, 64(3):511–533, 2012.
14. Jin-Yi Cai and Michael Kowalczyk. Spin systems on k-regular graphs with complex edge functions. *Theor. Comput. Sci.*, 461:2–16, 2012.
15. Jin-Yi Cai, Michael Kowalczyk, and Tyson Williams. Gadgets and anti-gadgets leading to a complexity dichotomy. In *ITCS*, pages 452–467. ACM, 2012.
16. Jin-Yi Cai and Pinyan Lu. Holographic algorithms with unsymmetric signatures. In *SODA*, pages 54–63. Society for Industrial and Applied Mathematics, 2008.
17. Jin-Yi Cai and Pinyan Lu. Holographic algorithms: From art to science. *J. Comput. Syst. Sci.*, 77(1):41–61, 2011.
18. Jin-Yi Cai and Pinyan Lu. Signature theory in holographic algorithms. *Algorithmica*, 61(4):779–816, 2011.
19. Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Holographic algorithms with matchgates capture precisely tractable planar #CSP. In *FOCS*, pages 427–436. IEEE Computer Society, 2010.
20. Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Computational complexity of Holant problems. *SIAM J. Comput.*, 40(4):1101–1132, 2011.
21. Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Dichotomy for Holant* problems of Boolean domain. In *SODA*, pages 1714–1728. SIAM, 2011.

22. Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Holographic reduction, interpolation and hardness. *Computational Complexity*, 21(4):573–604, 2012.
23. Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Dichotomy for Holant* problems with domain size 3. In *SODA*, pages 1278–1295. SIAM, 2013.
24. Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Holographic algorithms by Fibonacci gates. *Linear Algebra and its Applications*, 438(2):690–707, 2013.
25. Jin-Yi Cai, Pinyan Lu, and Mingji Xia. The complexity of complex weighted Boolean #CSP. *J. Comput. System Sci.*, 80(1):217–236, 2014.
26. Martin Dyer, Leslie Ann Goldberg, and Mark Jerrum. The complexity of weighted Boolean #CSP. *SIAM J. Comput.*, 38(5):1970–1986, 2009.
27. Martin Dyer, Leslie Ann Goldberg, and Mike Paterson. On counting homomorphisms to directed acyclic graphs. *J. ACM*, 54(6), 2007.
28. Martin Dyer and Catherine Greenhill. The complexity of counting graph homomorphisms. *Random Struct. Algorithms*, 17(3-4):260–289, 2000.
29. Martin Dyer and David Richerby. An effective dichotomy for the counting constraint satisfaction problem. *SIAM J. Comput.*, 42(3):1245–1274, 2013.
30. Leslie Ann Goldberg, Martin Grohe, Mark Jerrum, and Marc Thurley. A complexity dichotomy for partition functions with mixed signs. *SIAM J. Comput.*, 39(7):3336–3402, 2010.
31. Heng Guo, Sangxia Huang, Pinyan Lu, and Mingji Xia. The complexity of weighted Boolean #CSP modulo k. In *STACS*, pages 249–260. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2011.
32. Heng Guo and Tyson Williams. The complexity of planar Boolean #CSP with complex weights. In *ICALP*, pages 516–527. Springer Berlin Heidelberg, 2013.
33. Pavol Hell and Jaroslav Nešetřil. On the complexity of H-coloring. *J. Comb. Theory Ser. B*, 48(1):92–110, 1990.
34. Sangxia Huang and Pinyan Lu. A dichotomy for real weighted Holant problems. In *IEEE Conference on Computational Complexity*, pages 96–106. IEEE Computer Society, 2012.
35. P. W. Kasteleyn. The statistics of dimers on a lattice. *Physica*, 27(12):1209–1225, 1961.
36. P. W. Kasteleyn. Graph theory and crystal physics. In F. Harary, editor, *Graph Theory and Theoretical Physics*, pages 43–110. Academic Press, London, 1967.
37. Neeraj Kayal. Affine projections of polynomials: Extended abstract. In *STOC*, pages 643–662. ACM, 2012.
38. László Lovász. Operations with structures. *Acta Math. Hung.*, 18(3-4):321–328, 1967.
39. Ketan D. Mulmuley and Milind Sohoni. Geometric complexity theory I: An approach to the P vs. NP and related problems. *SIAM J. Comput.*, 31(2):496–526, 2001.
40. H. N. V. Temperley and Michael E. Fisher. Dimer problem in statistical mechanics—an exact result. *Philosophical Magazine*, 6(68):1061–1063, 1961.
41. Leslie G. Valiant. Expressiveness of matchgates. *Theor. Comput. Sci.*, 289(1):457–471, 2002.
42. Leslie G. Valiant. Quantum circuits that can be simulated classically in polynomial time. *SIAM J. Comput.*, 31(4):1229–1254, 2002.
43. Leslie G. Valiant. Accidental algorthims. In *FOCS*, pages 509–517. IEEE Computer Society, 2006.
44. Leslie G. Valiant. Holographic algorithms. *SIAM J. Comput.*, 37(5):1565–1594, 2008.