FPTAS for Weighted Fibonacci Gates and Its Applications

Pinyan Lu[1], Menghui Wang[2,⋆], and Chihao Zhang[3,⋆⋆]

[1] Microsoft Research
pinyanl@microsoft.com
[2] University of Wisconsin-Madison
menghui@cs.wisc.edu
[3] Shanghai Jiao Tong University
chihao.zhang@gmail.com

Abstract. Fibonacci gate problems have severed as computation primitives to solve other problems by holographic algorithm [5] and play an important role in the dichotomy of exact counting for Holant and CSP frameworks [6]. We generalize them to weighted cases and allow each vertex function to have different parameters, which is a much boarder family and #P-hard for exactly counting. We design a fully polynomial-time approximation scheme (FPTAS) for this generalization by correlation decay technique. This is the first deterministic FPTAS for approximate counting in the general Holant framework without a degree bound. We also formally introduce holographic reduction in the study of approximate counting and these weighted Fibonacci gate problems serve as computation primitives for approximate counting. Under holographic reduction, we obtain FPTAS for other Holant problems and spin problems. One important application is developing an FPTAS for a large range of ferromagnetic two-state spin systems. This is the first deterministic FPTAS in the ferromagnetic range for two-state spin systems without a degree bound. Besides these algorithms, we also develop several new tools and techniques to establish the correlation decay property, which are applicable in other problems.

1 Introduction

Holant is a refined framework for counting problems [5,6,8], which is more expressive than previous frameworks such as counting constraint satisfaction problems (CSP) in the sense that they can be simulated using Holant instances. In this paper, we consider a generalization called weighted Holant problems. A weighted Holant is an extension of a Holant problem where each edge e is assigned an activity λ_e, and if it is chosen it contributes to the partition function a factor of λ_e.

⋆ This work was partially performed when the author was an undergraduate student at Shanghai Jiao Tong University.

⋆⋆ The author is supported by NSF of China (61033002, ANR 61261130589)

⋆⋆⋆ The full version of this paper is available at http://arxiv.org/abs/1402.4370.

J. Esparza et al. (Eds.): ICALP 2014, Part I, LNCS 8572, pp. 787–799, 2014.

Given a graph $G(V, E)$, a family of node functions $\mathcal{F} = \{F_v | v \in V\}$, and edge weights $\Lambda = \{\lambda_e | e \in E\}$, the partition function for a weighted Holant instance $\Omega(G, \mathcal{F}, \Lambda)$ is the summation of the weights over all configurations $\sigma : E \to \{0, 1\}$, specifically the value of $\sum_\sigma \left(\prod_{e \in E} \lambda_e(\sigma(e)) \prod_{v \in V} F_v\left(\sigma|_{E(v)}\right)\right)$. We use Holant$(\mathcal{F}, \Lambda)$ to denote the class of Holant problems where all functions are taken from $\mathcal{F}$ and all edge weights are taken from Λ. For example, consider the Perfect Matching problem on G. This problem corresponds to attaching the Exact-One function on every vertex of G — for each 0-1 edge assignment, the product $\prod_{v \in V} F_v(\sigma |_{E(v)})$ evaluates to 1 when the assignment is a perfect matching, and 0 otherwise, thereby summing over all 0-1 edge assignments gives us the number of perfect matchings in G. If we use the At-Most-One function at each vertex, then we can count all matchings, including those that are not perfect.

A symmetric function F can be expressed by $[f_0, f_1, \ldots, f_k]$, where f_i is the value of F on inputs of hamming weight i. The above mentioned Exact-One and At-Most-One functions are both symmetric and can be expressed as $[0, 1, 0, 0, \ldots]$ and $[1, 1, 0, 0, \ldots]$ respectively. A Fibonacci function F is a symmetric function $[f_0, f_1, \ldots, f_k]$, satisfying that $f_i = cf_{i-1} + f_{i-2}$ for some constant c. For example, the parity function $[a, b, a, b, \ldots]$ is a special Fibonacci function with $c = 0$. If there are no edge weights (or equivalently all the weights are equal to 1) and all the node functions are Fibonacci functions with a same parameter c, we have a polynomial time algorithm to compute the partition function exactly [5]. These problems also form the base for a family of holographic algorithms, where other interesting problems can be reduced to the Fibonacci gate problems [5]. Furthermore, this family of functions is interesting not only because of its tractability, but also because it essentially captures almost all tractable Holant problems with all unary functions available [6,8].

If we allow edges to have non-trivial weights or each function to have different parameters in Fibonacci gates, then the exact counting problem becomes #P-hard [6,8]. Nevertheless, it is interesting to study the problem in the approximation setting. In contrast to the exact counting setting, the approximability of Holant problem is much less understood. In this paper, we study approximate counting for weighted Fibonacci gate problems.

Another closely related and well-studied model is spin systems. In this paper, we focus on two-state spin systems. An instance of a spin system is a graph $G(V, E)$. A configuration $\sigma : V \to \{0, 1\}$ assigns every vertex one of the two states. The contributions of local interactions between adjacent vertices are quantified by a matrix $A = \begin{bmatrix} A_{0,0} & A_{0,1} \\ A_{1,0} & A_{1,1} \end{bmatrix} = \begin{bmatrix} \beta & 1 \\ 1 & \gamma \end{bmatrix}$, where $\beta, \gamma \geq 0$. The partition function is defined by $Z_A(G) = \sum_{\sigma \in \{0,1\}^V} \prod_{(u,v) \in E} A_{\sigma(u),\sigma(v)}$.

There has been a lot of studies on the approximability of the partition function in terms of parameters β and γ. The problem is exactly solvable in polynomial time if $\beta\gamma = 1$. When $\beta\gamma < 1$, the system is called anti-ferromagnetic and we have a complete understanding of its approximability: there is a uniqueness boundary, above which there is an FPTAS [27,15,23,16] and below which it is NP-hard [24,25,9].

The story is different in ferromagnetic range $\beta\gamma > 1$. Jerrum and Sinclair [13] gave an FPRAS for Ising model ($\beta = \gamma > 1$) based on Markov Chain Monte Carlo (MCMC) method and lately Goldberg et al. extended that to all $\beta\gamma > 1$ plane. However, these algorithms are all randomized. Can we design a deterministic FPTAS for it as that for anti-ferromagnetic range? Indeed, this is an interesting and important question in general and many effort has been made for derandomizing MCMC based algorithms. For instance, there is an FPRAS for counting matchings [12] but FPTAS is only known for graphs of bounded degree [2]. The situation is similar in computing permanent of nonnegative matrix, although an FPRAS is known [14], the current best deterministic algorithm can only approximate the permanent with an exponential large factor [18]. To the best of our knowledge, no deterministic FPTAS was previously known for two-state spin systems in ferromagnetic range. In particular, the correlation decay technique, the main tool to design FPTAS in anti-ferromagnetic range, cannot directly apply.

1.1 Our Results

The main results of this paper are a number of FPTAS's for computing the partition function of different Holant problems and spin systems.

Weighted Fibonacci Gates. We design an FPTAS for weighted Fibonacci gates when the parameters satisfy certain conditions. We have several theorems to cover different ranges. In Theorem 1, we prove that for any fixed choice of other parameters, we can design an FPTAS as long as the edge weights are close enough to 1. This result demonstrates a smooth transition from the unweighted case to weighted ones in terms of approximation.

Another interesting range is that we have an FPTAS for the whole range as long as the Fibonacci parameter c is reasonably large (no less than a constant 1.17) and edge weights are no less than 1 (which means all the edges prefer to be chosen) (Theorem 2). We also allow different nodes to have functions with different parameter c, which contrasts the exact counting setting where a uniform parameter on each node is crucial to have a polynomial time algorithm.

Ferromagnetic Two-State Spin Systems. We design an FPTAS for a large range of ferromagnetic two state spin systems. This is the first deterministic FPTAS in the ferromagnetic range for two-state spin systems without a degree bound. To describe the tractable range, we present a monotonically increasing function $\Gamma : [1, \infty] \to \mathbb{R}$ with $\Gamma(1) = 1$ and $\Gamma(x) \le x$. We have an FPTAS for a ferromagnetic spin system $\begin{bmatrix} \beta & 1 \\ 1 & \gamma \end{bmatrix}$ as long as $\gamma \le \Gamma(\beta)$ or $\beta \le \Gamma(\gamma)$ (Theorem 4). The exact formula of Γ is complicated and we do not spend much effort to optimize it. However, it already enjoys a nice property in that $\lim_{x \to +\infty} \frac{\Gamma(x)}{x} = 1$. This means that although the range does not cover the Ising model ($\beta = \gamma$), it gets relatively close to that in infinity. We also have similar results for two-spin system with external fields.

Other Holant Problems. We can extend our FPTAS to functions $[f_0, f_1, \ldots, f_d]$ with form $f_{i+2} = af_{i+1} + bf_i$ for a range of parameters. This is a much bigger family than Fibonacci gates, since Fibonacci gates corresponds to $b = 1$.

1.2 Our Techniques

Our main approach for designing FPTAS's is the correlation decay technique introduced in [1] and [27]. While the general framework is standard, it is highly non-trivial to design a recursive computational structure and especially to prove the property of exponential correlation decay for a specific problem.

A powerful technique we use is to apply a potential function to amortize the decay rate, which was introduced and used in many circumstances [22,15,23,16,20]. Besides this, to enrich the toolkit, we introduce several new techniques to design and analysis the recursive computational structure. We believe that these techniques can find their applications in other problems.

Working with Dangling Edges. The recursive computational structure for spin problems usually relates a marginal probability of a vertex to that of its neighbors. In Holant problems, we work with assignments and marginal probabilities on edges. Since an edge has two ends, it has two sets of neighbors, which complicates things a lot. In this paper, we instead work on instances with dangling edges, that is, a half edge with neighbors only on one end, and then reduce regular instances to dangling instances. This technique works for any Holant problems and we believe that it is the right structure to work with in the Holant framework. Indeed, the idea has later been successfully used in [17].

Computation Tree with Bounded Degrees. The correlation decay property only directly implies an FPTAS for systems with bounded degrees. One exception is the anti-ferromagnetic two-state spin systems, where a stronger notion of computationally efficient correlation decay is introduced [15]. In this paper, we also establish the computationally efficient correlation decay for systems with unbounded degree, but via a different approach. Thanks to the unique property of Fibonacci functions, we can decompose a node into several nodes with constant degrees. Thus, at each step of our computation tree, we only involve constant many sub-instances even if the degree of the original system is not bounded.

Bounding Range of Variables. After we get a recursion system, the main task is to prove the correlation decay property. This is usually achieved by proving that a certain amortized decay rate, which is a function of several variables, is less than one for any choice of these variables in their domain. Thus if one can prove a smaller domain for each vairable, the analysis becomes easier. Some naive implementation of this idea already appeared in approximate counting of coloring problems [10,20]. In this paper, we develop this idea much further. We divide sub-instances involved in the computation tree into two classes: deep ones for which we can get a much better estimation of their range and shallow ones for which we can compute their value without error. Then we can either compute the exact value or we can safely assume that it is within a smaller domain, which enables us to prove the correlation decay property easier.

Holographic Reduction. We formally introduce holographic reduction in the study of approximate counting. We use weighted Fibonacci gate problems as computational primitives for approximate counting and design holographic algorithms for other problems based on them. In particular, we use the FPTAS for Fibonacci gates to obtain an FPTAS for ferromagnetic two-state spin systems. It is noteworthy that the correlation decay property does not generally hold for ferromagnetic two-state spin systems. So we cannot do a similar argument to get the FPTAS in the spin world directly. Moreover, the idea of holographic reduction can apply to any Holant problems, which extends known counting algorithms (both exact and approximate, both deterministic and randomized) to a broader family of problems. Indeed, the other direction of holographic reduction is also used in our algorithm. We design an exact algorithm for shallow sub-instances of Fibonacci instance by a holographic reduction to the spin world.

1.3 Related Works

Most previous studies of the Holant framework are for exact counting, and a number of dichotomy theorems were proved [8,11,3]. Holographic reduction was introduced by Valiant in holographic algorithms [26,4], which is later also used to prove hardness result of counting problems [5,8,7].

For some special Holant problems such as counting (perfect) matchings, their approximate versions are well studied [2,12,14]. In particular, [2] gave an FPTAS to count matchings but only for graphs with bounded degrees. It is relatively less studied in the general Holant framework in terms of approximate counting except for two recent work: [28] studied general Holant problems but only for planar graph instances with a bounded degree; [21] gives an FPRAS for several Holant problems. Another well-known example is the "sub-graph world" in [13]. It is indeed a weighted Holant problem with Fibonacci functions of $c = 0$, for which an FPRAS was given. In that paper, holographic reduction was also implicitly used, which extends the FPRAS to the Ising model.

Most previous study for FPTAS via correlation decay is on the spin systems. It was extremely successful in the anti-ferromagnetic two-spin system [27,15,23,16]. It is also used in multi-spin systems [10,20].

2 Preliminaries

A weighted Holant instance $\Omega = (G(V,E), \{F_v|v \in V\}, \{\lambda_e|e \in E\})$ is a tuple. $G(V,E)$ is a graph. F_v is a function with arity d_v: $\{0,1\}^{d_v} \to \mathbb{R}^+$, where d_v is the degree of v and $\mathbb{R}^+$ denotes non-negative real numbers. Edge weight λ_e is a mapping $\{0,1\} \to \mathbb{R}^+$. A configuration σ is a mapping $E \to \{0,1\}$ and gives a weight $w_\Omega(\sigma) = \prod_{e\in E} \lambda_e(\sigma(e)) \prod_{v\in V} F_v(\sigma\mid_{E(v)})$, where $E(v)$ denotes the incident edges of v. The counting problem on the instance Ω is to compute the partition function: $Z(\Omega) = \sum_\sigma \left(\prod_{e\in E} \lambda_e(\sigma(e)) \prod_{v\in V} F_v(\sigma\mid_{E(v)})\right)$.

We can represent each function F_v by a vector in $(\mathbb{R}^+)^{2^{d_v}}$, or a tensor in $((\mathbb{R}^+)^2)^{\otimes d_v}$. This is also called a *signature*. A symmetric function F can be expressed by $[f_0, f_1, \ldots, f_k]$, where f_j is the value of F on inputs of hamming weight j. For example, the equality function is $[1, 0, \ldots, 0, 1]$. Edge weight is a unary function, which can be written as $[\lambda_e(0), \lambda_e(1)]$. Since we do not care about global scale factor, we always normalize that $\lambda_e(0) = 1$ and use the notation $\lambda_e = \lambda_e(1)$ as a real number.

A Holant problem is parameterized by a set of functions $\mathcal{F}$ and edge weights Λ. We denote by Holant$(\mathcal{F}, \Lambda)$ the following computation problem .

Definition 1. *Given a set of functions $\mathcal{F}$ and edge weights Λ, we denote by* Holant$(\mathcal{F}, \Lambda)$ *the following computation problem.*
Input: *A Holant instance $\Omega = (G(V, E), \{F_v | v \in V\}, \{\lambda_e | e \in E\})$, where $F_v \in \mathcal{F}$ and $\lambda_e \in \Lambda$;*
Output: *The partition function $Z(\Omega)$.*

The weights of configurations also give a distribution over all possible configurations:

$$\mathbb{P}_\Omega(\sigma) = \frac{w_\Omega(\sigma)}{Z(\Omega)} = \frac{1}{Z(\Omega)} \prod_{e \in E} \lambda_e(\sigma(e)) \prod_{v \in V} F_v(\sigma |_{E(v)}).$$

This defines the marginal probability of each edge $e_0 \in E$.

$$\mathbb{P}_\Omega(\sigma(e_0) = 0) = \frac{\sum_{\sigma: \sigma(e_0)=0} \left(\prod_{e \in E} \lambda_e(\sigma(e)) \prod_{v \in V} F_v(\sigma |_{E(v)})\right)}{Z(\Omega)}.$$

Similarly, we can define the marginal probability of a subset of edges. Let $E_0 \subset E$ and $e_1, e_2, \ldots, e_{|E_0|}$ be an enumeration of the edges in E_0. Then we can define $\sigma(E_0) = \sigma(e_1)\sigma(e_2) \cdots \sigma(e_{|E_0|})$ as a Boolean string of length $|E_0|$. Let $\alpha \in \{0, 1\}^{|E_0|}$, we define

$$\mathbb{P}_\Omega(\sigma(E_0) = \alpha) = \frac{\sum_{\sigma: \sigma(e_i)=\alpha_i, i=1,2,\ldots,|E_0|} \left(\prod_{e \in E} \lambda_e(\sigma(e)) \prod_{v \in V} F_v(\sigma |_{E(v)})\right)}{Z(\Omega)}.$$

We denote the partial summation as

$$Z(\Omega, \sigma(E_0) = \alpha) = \sum_{\sigma: \sigma(e_i)=\alpha_i} \left(\prod_{e \in E} \lambda_e(\sigma(e)) \prod_{v \in V} F_v(\sigma |_{E(v)}) \right).$$

We define a dangling instance Ω^D of Holant$(\mathcal{F}, \Lambda)$ also as a tuple $(G(V, E \cup D), \{F_v | v \in V\}, \{\lambda_e | e \in E\})$, where $G(V, E \cup D)$ is a graph with dangling edges D. A dangling edge can be viewed as a half edge, with one end attached to a regular vertex in V and the other end dangling (not considered as a vertex). A dangling instance Ω^D is the same as a Holant instance except for these dangling edges. In $G(V, E \cup D)$ 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), each regular edge in E is assigned a weight from Λ and we always assume that there is no weight on a dangling edge in this paper. A dangling instance can be also viewed as a regular instance by attaching a vertex with function $[1, 1]$ at the dangling end of each dangling edge. We can define the probability distribution and marginal probabilities just as for regular instance. In particular, we shall use dangling instance Ω^e with single dangling edge e extensively in this paper. For that, we define $R(\Omega^e) = \frac{\mathbb{P}_{\Omega^e}(\sigma(e)=1)}{\mathbb{P}_{\Omega^e}(\sigma(e)=0)}$.

3 Statement of Main Results

A symmetrical function $[f_0, f_1, \ldots, f_d]$ is called a (generalized) Fibonacci function if there exists a constant c such that $f_{i+2} = cf_{i+1} + f_i$, where $i = 0, 1, \cdots, d-2$. We denote this family of function as $\mathcal{F}_c$, the Fibonacci functions with parameter c. We use $\mathcal{F}_c^{p,q}$ to denote a subfamily of $\mathcal{F}_c$ such that $f_{i+1} \geq pf_i$ and $f_{i+1} \leq qf_i$ for all $i = 0, 1, \cdots, d-1$. When the upper bound q is not given, we simply write $\mathcal{F}_c^p$. We use $\mathcal{F}_{c_1,c_2}^{p,q}$ to denote $\bigcup_{c_1 \leq c \leq c_2} \mathcal{F}_c^{p,q}$. We use $\Lambda_{\lambda_1,\lambda_2}$ to denote the set of edge weights λ_e such that $\lambda_1 \leq \lambda_e \leq \lambda_2$.

Here is a list of FPTAS's we get:

Theorem 1. *For any $c > 0$ and $p > 0$, there exists $\lambda_1(p, c) < 1$ and $\lambda_2(p, c) > 1$ such that there is an FPTAS for* $\text{Holant}(\mathcal{F}_c^p, \Lambda_{\lambda_1(p,c),\lambda_2(p,c)})$.

Theorem 2. *Let $p > 0$. Then there is an FPTAS for* $\text{Holant}(\mathcal{F}_{1.17,+\infty}^p, \Lambda_{1,+\infty})$.

Theorem 3. *Let $\lambda > 0$ and $c \geq 2.57$. There is an FPTAS for* $\text{Holant}(\mathcal{F}_c^{c/2,c+2/c}, \Lambda_{\lambda,+\infty})$.

Under a holographic reduction with base $\begin{bmatrix} 1 & t \\ \rho & -\frac{t}{\rho} \end{bmatrix}$, we have the following transformation. Let $\lambda > 0$, $\rho \geq 1$, $t(1-\lambda) > 0$, and $|t| \leq 1$. Let $\beta = \frac{1+\lambda\rho^2}{t(1-\lambda)}$ and $\gamma = \frac{t(1+\lambda\rho^{-2})}{1-\lambda}$. The two spin problem with edge function $\begin{bmatrix} \beta & 1 \\ 1 & \gamma \end{bmatrix}$ and external field μ is equivalent to $\text{Holant}(\mathcal{F}_{\rho-\frac{1}{\rho}}, \Lambda_{\lambda,\lambda})$, where $\mathcal{F}_{\rho-\frac{1}{\rho}}$ is a set of Fibonacci functions with with parameter $c = \rho - \frac{1}{\rho}$ and the one of arity n has form $f_k = \rho^k + \mu t^n(-\rho)^{-k}$. Through this reduction, we can transform Theorem 1-3 to the following FPTAS for ferromagnetic two-state spin system .

Theorem 4. *There is a continuous curve $\Gamma(\beta)$ defined on $[1, +\infty)$ such that (1) $\Gamma(1) = 1$; (2) $1 < \Gamma(\beta) < \beta$ for all $\beta > 1$; and (3) $\lim_{\beta\to+\infty} \frac{\Gamma(\beta)}{\beta} = 1$. There is an FPTAS for the two-state spin system with local interaction matrix $\begin{bmatrix} \beta & 1 \\ 1 & \gamma \end{bmatrix}$ and external field $\mu \leq 1$ if $\beta\gamma > 1$ and $\gamma \leq \Gamma(\beta)$.*

4 Computation Tree Recursion

In the exact polynomial time algorithm for Fibonacci gates without edge weights, one crucial property of a set of Fibonacci functions with a fixed parameter is

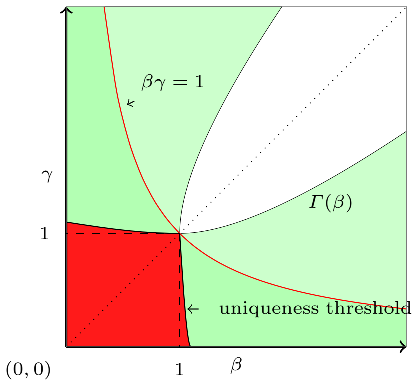

Fig. 1. This figure illustrates the rough shape of $\Gamma(\cdot)$ when there is no external field. It also includes anti-ferromagnetic range. Parameters (β, γ) admit FPTAS in green region and hard to approximate in red region.

that it is closed when two nodes are connected together [5]. This is no longer true if we have non-trivial edge weights or when different Fibonacci function have different parameters. However, we can still use the special property of a Fibonacci function to decompose a vertex, which is the key property for all FPTAS algorithms in our paper.

Let $\Omega = (G(V,E), \{F_v | v \in V\}, \{\lambda_e | e \in E\})$ be an instance of Holant $(\mathcal{F}^{p,q}_{c_1,c_2}, \Lambda_{\lambda_1,\lambda_2})$, $v \in V$ be a vertex of the instance with degree $d_1 + d_2$ ($d_1, d_2 \geq 1$) and $e_1, e_2, \ldots, e_{d_1+d_2}$ be its incident edges. We can construct a new Holant instance Ω': Ω' is the same as Ω except that v is decomposed into two vertices v', v''. $e_1, e_2, \ldots, e_{d_1}$ are connected to v' and $e_{d_1+1}, e_{d_1+2}, \ldots, e_{d_1+d_2}$ are connected to v''. There is a new edge e connecting v' and v''. If the function on the original v is $[f_0, f_1, \ldots, f_{d_1+d_2}]$, a Fibonacci function with parameter c, then the function on v' is $[f_0, f_1, \ldots, f_{d_1}]$ and the function on v'' is $[1, 0, 1, c \ldots]$, also a Fibonacci function with parameter c. The edge weight on the new edge e is 1. The functions on all other nodes and edge weights on all other edges (except the new e) remain the same as that in Ω. We use the following notation to denote this decomposition operation

$$\Omega' = D(\Omega, v, \{e_1, e_2, \ldots, e_{d_1}\}, \{e_{d_1+1}, e_{d_1+2}, \ldots, e_{d_1+d_2}\}).$$

Using the special property of Fibonacci function, we have the following lemma.

Lemma 1. *Let $\Omega' = D(\Omega, v, E_1, E_2)$. Then $Z(\Omega) = Z(\Omega')$ and for all $e \in E$, $\mathbb{P}_{\Omega}(\sigma(e) = 0) = \mathbb{P}_{\Omega'}(\sigma(e) = 0)$.*

Let Ω^e be a dangling instance of Holant$(\mathcal{F}^{p}_{c_1,c_2}, \Lambda_{\lambda_1,\lambda_2})$. Let v be the attaching vertex of the dangling edge e and $e_1, e_2, \ldots, e_d$ be other incident edges of v. We compute $R(\Omega^e)$ by smaller instances depending on d. If $d = 0$, then $R(\Omega^e)$ can be computed directly. If $d = 1$, we construct a smaller dangling instance Ω^{e_1} by removing e_0 and v from G and make e_1 be the new dangling edge and remove its weight.

$$R(\Omega^e) = \frac{f_1 + \lambda_{e_1} f_2 R(\Omega^{e_1})}{f_0 + \lambda_{e_1} f_1 R(\Omega^{e_1})}. \tag{1}$$

If $d \geq 2$, we use the above lemma to decompose the vertex v into v' and v'' and let e and e_1 connect to v'' and the remaining edges connect to v'. We use e' to denote the edge between v' and v''. By removing e and v'' from Ω' , we get a dangling instance Ω^{e',e_1} with two dangling edges e', e_1.

$$\begin{aligned}
R(\Omega^e) &= \frac{Z(\Omega^e, \sigma(e)=1)}{Z(\Omega^e, \sigma(e)=0)} \\
&= \frac{\lambda_{e_1} Z(\Omega^{e',e_1}, \sigma(e'e_1)=01) + Z(\Omega^{e',e_1}, \sigma(e'e_1)=10) + c\lambda_{e_1} Z(\Omega^{e',e_1}, \sigma(e'e_1)=11)}{Z(\Omega^{e',e_1}, \sigma(e'e_1)=00) + \lambda_{e_1} Z(\Omega^{e',e_1}, \sigma(e'e_1)=11)} \\
&= \frac{\lambda_{e_1} \mathbb{P}_{\Omega^{e',e_1}}(\sigma(e'e_1)=01) + \mathbb{P}_{\Omega^{e',e_1}}(\sigma(e'e_1)=10) + c\lambda_{e_1} \mathbb{P}_{\Omega^{e',e_1}}(\sigma(e'e_1)=11)}{\mathbb{P}_{\Omega^{e',e_1}}(\sigma(e'e_1)=00) + \lambda_{e_1} \mathbb{P}_{\Omega^{e',e_1}}(\sigma(e'e_1)=11)}.
\end{aligned}$$

In the above recursion, the marginal probability of the original instance is written as that of smaller instances but with two dangling edges. In order to continue the recursive process, we need to convert them into instances with single dangling edge. This can be done by pinning one of the two dangling edges, or just leaving one of the edges free (in which case the dangling end of the free edge can be treated as a regular vertex with signature $[1, 1]$).

We use $\textsc{Pin}_{e,x}(\Omega)$ to denote the new instance obtained by pinning the edge e of the instance Ω to x.

There are many choices in deciding which edge to pin, and to what state the edge is pinned to. Each choice leads to different recursions and consequently have an impact on the following analysis. Here we give an example which is used in the proof of Theorem 1 and Theorem 3. In the proof of Theorem 2, we use a different one.

Set $\Omega^{e'} = \textsc{Pin}_{e_1,0}(\Omega^{e',e_1})$, $\Omega^{e_1} = \textsc{Pin}_{e',0}(\Omega^{e',e_1})$ and $\widetilde{\Omega}^{e_1} = \textsc{Pin}_{e',1}(\Omega^{e',e_1})$. By the definitions, we have $\mathbb{P}_{\Omega^{e'}}(\sigma(e') = 0) = \mathbb{P}_{\Omega^{e',e_1}}(\sigma(e') = 0|\sigma(e_1) = 0)$, $\mathbb{P}_{\Omega^{e_1}}(\sigma(e_1) = 0) = \mathbb{P}_{\Omega^{e',e_1}}(\sigma(e_1) = 0|\sigma(e') = 0)$, and $\mathbb{P}_{\widetilde{\Omega}^{e_1}}(\sigma(e_1) = 0) = \mathbb{P}_{\Omega^{e',e_1}}(\sigma(e_1) = 0|\sigma(e') = 1)$. Given these relation and the fact that

$$\mathbb{P}_{\Omega^{e',e_1}}(\sigma(e'e_1) = 00) + \mathbb{P}_{\Omega^{e',e_1}}(\sigma(e'e_1) = 01) +$$

$$\mathbb{P}_{\Omega^{e',e_1}}(\sigma(e'e_1) = 10) + \mathbb{P}_{\Omega^{e',e_1}}(\sigma(e'e_1) = 11) = 1.$$

We can solve these marginal probabilities and substitute these into the above recursion to get

$$R(\Omega^e) = \frac{\lambda_{e_1} R(\Omega^{e_1}) + R(\Omega^{e'}) + c\lambda_{e_1} R(\Omega^{e'}) R(\widetilde{\Omega}^{e_1})}{1 + \lambda_{e_1} R(\Omega^{e'}) R(\widetilde{\Omega}^{e_1})} \tag{2}$$

If e' and e_1 are in different connected components of Ω^{e',e_1}, then the marginal probability of e_1 is independent of e' and as a result $R(\widetilde{\Omega}^{e_1}) = R(\Omega^{e_1})$. So in this case, we have

$$R(\Omega^e) = \frac{\lambda_{e_1} R(\Omega^{e_1}) + R(\Omega^{e'}) + c\lambda_{e_1} R(\Omega^{e'})R(\Omega^{e_1})}{1 + \lambda_{e_1} R(\Omega^{e'})R(\Omega^{e_1})} \tag{3}$$

Starting from an dangling instance Ω^e, we can compute $R(\Omega^e)$ by one of (1), (2) and (3) recursively. We note that if $\Omega^e \in \text{Holant}(\mathcal{F}^{p,q}_{c_1,c_2}, \Lambda_{\lambda_1,\lambda_2})$, the instances involved in the recursion are also in the same family. We define three functions according to these three recursions:

$$h(x) = \frac{f_1 + \lambda_{e_1} f_2 x}{f_0 + \lambda_{e_1} f_1 x}, g(x,y,z) = \frac{\lambda_{e_1} y + x + c\lambda_{e_1} xz}{1 + \lambda_{e_1} xz}, \hat{g}(x,y) = \frac{\lambda_{e_1} y + x + c\lambda_{e_1} xy}{1 + \lambda_{e_1} xy}.$$

By expanding this recursion, we get a computation tree recursion to compute $R(\Omega^e)$. We need one more step to compute the marginal probability of an edge in a regular instance. This can be done similarly and we have the following lemma.

Lemma 2. *If we can ϵ approximate $R(\Omega^e)$ for any dangling instance Ω^e of* $\text{Holant}(\mathcal{F}^{p,q}_{c_1,c_2}, \Lambda_{\lambda_1,\lambda_2})$ *in time poly$(n, \frac{1}{\epsilon})$, we can also ϵ approximate the marginal probability of any edge of a regular instance of* $\text{Holant}(\mathcal{F}^{p,q}_{c_1,c_2}, \Lambda_{\lambda_1,\lambda_2})$ *in time poly$(n, \frac{1}{\epsilon})$.*

5 Algorithm

The procedure from marginal probabilities to partition function is rather standard and we have the following lemma.

Lemma 3. *If for any $\epsilon > 0$ and any Ω^e of* $\text{Holant}\left(\mathcal{F}^{p,q}_{c_1,c_2}, \Lambda_{\lambda_1,\lambda_2}\right)$*, we have a deterministic algorithm to get $\widehat{P}$ in time poly $\left(n, \frac{1}{\epsilon}\right)$ such that $|\widehat{P} - \mathbb{P}_{\Omega^e}(\sigma(e) = 0)| \le \epsilon$, we have an FPTAS for* $\text{Holant}(\mathcal{F}^{p,q}_{c_1,c_2}, \Lambda_{\lambda_1,\lambda_2})$.

Before we use the computation tree recursion to compute the marginal probability, we need the following lemma to handle shallow instances separately. We denote by $SP(\Omega^e)$ the longest simple path containing e in G.

Lemma 4. *Let L be a constant. We have a polynomial time algorithm to compute $R(\Omega^e)$ for all Ω^e of* $\text{Holant}(\mathcal{F}^{p}_{c_1,c_2}, \Lambda_{\lambda_1,\lambda_2})$ *with $SP(\Omega^e) \le L$.*

The proof of the above Lemma uses holographic reduction to spin world and makes use of the self-avoiding walk tree [27] for two-state spin systems. The length of the longest simple path is the same as the depth of the self-avoiding walk tree. See the full version for more details.

Now we give out formal procedure to estimate $\mathbb{P}_{\Omega^e}(\sigma(e) = 0)$. Since there is a one to one relation between $\mathbb{P}_{\Omega^e}(\sigma(e) = 0)$ and $R(\Omega^e)$, we can define our recursion on $R(\Omega^e)$, and at the final step we convert $R(\Omega^e)$ back to $\mathbb{P}_{\Omega}(\sigma(e) = 0)$. Let bounds R_1, R_2 and depth L be obtained for the family of dangling instance in the sense that for any dangling instance with $SP(\Omega^e) \geq L$, we have $R(\Omega^e) \in [R_1, R_2]$. Formally, for $t \geq 0$, the quantity $R^t(\Omega^e)$ is recursively defined as follows:

- If $SP(\Omega^e) \leq 2L$, we compute $R^t(\Omega^e) = R(\Omega^e)$ by Lemma 4.
- Else If $t = 0$, let $R^0(\Omega^e) = R_1$.
- Else If $t > 0$, use one of the recursion to get $\tilde{R}^t(\Omega^e) = \hat{g}(R^{t-1}(\Omega^{e'}), R^{t-1}(\Omega^{e_1}))$, $\tilde{R}^t(\Omega^e) = h(R^{t-1}(\Omega^{e_1}))$, or $\tilde{R}^t(\Omega^e) = g(R^{t-1}(\Omega^{e'}), R^{t-1}(\Omega^{e_1}), R^{t-1}(\widetilde{\Omega}^{e_1})$. Return the median of $R_1, \tilde{R}^t(\Omega^e), R_2$: $R^t(\Omega^e) = Med(R_1, \tilde{R}^t(\Omega^e), R_2)$.

There are three possible recursions and we define four amortized decay rates:

$$\alpha_1(x) = \frac{\Phi(x)\left|\frac{\mathrm{d}\,h}{\mathrm{d}\,x}\right|}{\Phi(h(x))}, \quad \alpha_3(x,y) = \frac{\left|\frac{\partial \hat{g}}{\partial x}\right|\Phi(x)}{\Phi(\hat{g}(x,y))}, \quad \alpha_4(x,y) = \frac{\left|\frac{\partial \hat{g}}{\partial y}\right|\Phi(y)}{\Phi(\hat{g}(x,y))},$$

$$\alpha_2(x,y,z) = \frac{1}{\Phi(g(x,y,z))}\left(\left|\frac{\partial g}{\partial x}\right|\Phi(x) + \left|\frac{\partial g}{\partial y}\right|\Phi(y) + \left|\frac{\partial g}{\partial z}\right|\Phi(z)\right),$$

where $\Phi(\cdot)$ is a potential function.

Definition 2. *We call a function* $\Phi : (0, +\infty) \to (0, +\infty)$ nice *if there is some function* $f : [1, +\infty) \to (0, +\infty)$ *such that for any* $c \geq 1$ *and* $x, y > 0$ *with* $\frac{x}{c} \leq y \leq cx$, *we have* $\frac{\Phi(x)}{\Phi(y)} \leq f(c)$.

Lemma 5. *Let bounds* R_1, R_2 *and depth* L *be obtained for dangling instances of* Holant$(\mathcal{F}^{p,q}_{c_1,c_2}, \Lambda_{\lambda_1,\lambda_2})$ *such that for any dangling instance with* $SP(\Omega^e) \geq L$, *we have* $R(\Omega^e) \in [R_1, R_2]$. *If there exist a nice function* $\Phi(\cdot)$ *and a constant* $\alpha < 1$ *such that* $\alpha_1(x) \leq \alpha$ *for all* $x \in [R_1, R_2]$, $\alpha_2(x,y,z) \leq \alpha$ *for all* $x, y, z \in [R_1, R_2]$, $\alpha_3(x,y) \leq \alpha$ *for all* $x \in [R_1, R_2]$, *and* $\alpha_4(x,y) \leq \alpha$ *for all* $y \in [R_1, R_2]$. *Then there is an FPTAS for* Holant$(\mathcal{F}^{p,q}_{c_1,c_2}, \Lambda_{\lambda_1,\lambda_2})$.

To obtain the FPTAS for the Fibonacci gates (Theorem 1-3), we make use of this Lemma 5. In order to apply Lemma 5, we need to establish two things: the bounds R_1, R_2 and the amortized decay rates. There two parts are technically involved and omitted here due to space limitation. The complete proof can be found in the full version of the current paper [19].

References

1. Bandyopadhyay, A., Gamarnik, D.: Counting without sampling: Asymptotics of the log-partition function for certain statistical physics models. Random Structures & Algorithms 33(4), 452–479 (2008)
2. Bayati, M., Gamarnik, D., Katz, D., Nair, C., Tetali, P.: Simple deterministic approximation algorithms for counting matchings. In: Proceedings of STOC, pp. 122–127. ACM (2007)

3. Cai, J.-Y., Guo, H., Williams, T.: A complete dichotomy rises from the capture of vanishing signatures. In: Proceedings of STOC, pp. 635–644. ACM (2013)
4. Cai, J.-Y., Lu, P.: Holographic algorithms: From art to science. Journal of Computer and System Sciences 77(1), 41–61 (2011)
5. Cai, J.-Y., Lu, P., Xia, M.: Holographic algorithms by Fibonacci gates and holographic reductions for hardness. In: Proceedings of FOCS, pp. 644–653 (2008)
6. Cai, J.-Y., Lu, P., Xia, M.: Holant problems and counting CSP. In: Proceedings of STOC, pp. 715–724 (2009)
7. Cai, J.-Y., Lu, P., Xia, M.: Holographic algorithms with matchgates capture precisely tractable planar #CSP. In: Proceedings of FOCS, pp. 427–436. IEEE (2010)
8. Cai, J.-Y., Lu, P., Xia, M.: Computational complexity of Holant problems. SIAM Journal on Computing 40(4), 1101–1132 (2011)
9. Galanis, A., Stefankovic, D., Vigoda, E.: Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. arXiv preprint arXiv:1203.2226 (2012)
10. Gamarnik, D., Katz, D.: Correlation decay and deterministic FPTAS for counting colorings of a graph. Journal of Discrete Algorithms 12, 29–47 (2012)
11. Huang, S., Lu, P.: A dichotomy for real weighted Holant problems. In: Proceedings of CCC, pp. 96–106. IEEE (2012)
12. Jerrum, M., Sinclair, A.: Approximating the permanent. SIAM Journal on Computing 18(6), 1149–1178 (1989)
13. Jerrum, M., Sinclair, A.: Polynomial-time approximation algorithms for the Ising model. SIAM Journal on Computing 22(5), 1087–1116 (1993)
14. Jerrum, M., Sinclair, A., Vigoda, E.: A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries. Journal of the ACM 51(4), 671–697 (2004)
15. Li, L., Lu, P., Yin, Y.: Approximate counting via correlation decay in spin systems. In: Proceedings of SODA, pp. 922–940. SIAM (2012)
16. Li, L., Lu, P., Yin, Y.: Correlation decay up to uniqueness in spin systems. In: Proceedings of SODA, pp. 67–84 (2013)
17. Lin, C., Liu, J., Lu, P.: A simple FPTAS for counting edge covers. In: Proceedings of SODA, pp. 341–348 (2014)
18. Linial, N., Samorodnitsky, A., Wigderson, A.: A deterministic strongly polynomial algorithm for matrix scaling and approximate permanents. Combinatorica 20(4), 545–568 (2000)
19. Lu, P., Wang, M., Zhang, C.: FPTAS for weighted Fibonacci gates and its applications. arXiv preprint arXiv:1402.4370 (2014)
20. Lu, P., Yin, Y.: Improved FPTAS for multi-spin systems. In: Raghavendra, P., Raskhodnikova, S., Jansen, K., Rolim, J.D.P. (eds.) RANDOM 2013 and APPROX 2013. LNCS, vol. 8096, pp. 639–654. Springer, Heidelberg (2013)
21. McQuillan, C.: Approximating Holant problems by winding. arXiv preprint arXiv:1301.2880 (2013)
22. Restrepo, R., Shin, J., Tetali, P., Vigoda, E., Yang, L.: Improved mixing condition on the grid for counting and sampling independent sets. Probability Theory and Related Fields 156(1-2), 75–99 (2013)
23. Sinclair, A., Srivastava, P., Thurley, M.: Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs. In: Proceedings of SODA, pp. 941–953. SIAM (2012)
24. Sly, A.: Computational transition at the uniqueness threshold. In: Proceedings of FOCS, pp. 287–296. IEEE (2010)

25. Sly, A., Sun, N.: The computational hardness of counting in two-spin models on d-regular graphs. In: Proceedings of FOCS, pp. 361–369. IEEE (2012)
26. Valiant, L.G.: Holographic algorithms. SIAM Journal on Computing 37(5), 1565–1594 (2008)
27. Weitz, D.: Counting independent sets up to the tree threshold. In: Proceedings of STOC, pp. 140–149. ACM (2006)
28. Yin, Y., Zhang, C.: Approximate counting via correlation decay on planar graphs. In: Proceedings of SODA, pp. 47–66. SIAM (2013)