

Inapproximability of the independent set polynomial in the complex plane

Ivona Bezáková* Andreas Galanis† Leslie Ann Goldberg† Daniel Štefankovič‡

November 2, 2017

Abstract

We study the complexity of approximating the value of the independent set polynomial $Z_G(\lambda)$ of a graph G with maximum degree Δ when the activity λ is a complex number.

When λ is real, the complexity picture is well-understood, and is captured by two real-valued thresholds λ^* and λ_c , which depend on Δ and satisfy $0 < \lambda^* < \lambda_c$. It is known that if λ is a real number in the interval $(-\lambda^*, \lambda_c)$ then there is an FPTAS for approximating $Z_G(\lambda)$ on graphs G with maximum degree at most Δ . On the other hand, if λ is a real number outside of the (closed) interval, then approximation is NP-hard. The key to establishing this picture was the interpretation of the thresholds λ^* and λ_c on the Δ -regular tree. The “occupation ratio” of a Δ -regular tree T is the contribution to $Z_T(\lambda)$ from independent sets containing the root of the tree, divided by $Z_T(\lambda)$ itself. This occupation ratio converges to a limit, as the height of the tree grows, if and only if $\lambda \in [-\lambda^*, \lambda_c]$.

Unsurprisingly, the case where λ is complex is more challenging. It is known that there is an FPTAS when λ is a complex number with norm at most λ^* and also when λ is in a small strip surrounding the real interval $[0, \lambda_c)$. However, neither of these results is believed to fully capture the truth about when approximation is possible. Peters and Regts identified the values of λ for which the occupation ratio of the Δ -regular tree converges. These values carve a cardioid-shaped region Λ_Δ in the complex plane, whose boundary includes the critical points $-\lambda^*$ and λ_c . Motivated by the picture in the real case, they asked whether Λ_Δ marks the true approximability threshold for general complex values λ .

Our main result shows that for every λ outside of Λ_Δ , the problem of approximating $Z_G(\lambda)$ on graphs G with maximum degree at most Δ is indeed NP-hard. In fact, when λ is outside of Λ_Δ and is not a positive real number, we give the stronger result that approximating $Z_G(\lambda)$ is actually #P-hard. Further, on the negative real axis, when $\lambda < -\lambda^*$, we show that it is #P-hard to even decide whether $Z_G(\lambda) > 0$, resolving in the affirmative a conjecture of Harvey, Srivastava and Vondrák.

Our proof techniques are based around tools from complex analysis — specifically the study of iterative multivariate rational maps. The full version is available at arxiv.org/abs/1711.00282 and is attached as an appendix. The theorem numbering here matches the full version.

*Department of Computer Science, Rochester Institute of Technology, Rochester, NY, USA. Research supported by NSF grant CCF-1319987.

†Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, Oxford, OX1 3QD, UK. The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) ERC grant agreement no. 334828. The paper reflects only the authors’ views and not the views of the ERC or the European Commission. The European Union is not liable for any use that may be made of the information contained therein.

‡Department of Computer Science, University of Rochester, Rochester, NY 14627. Research supported by NSF grant CCF-0910415.

1 Introduction

The independent set polynomial is one of the most well-studied graph polynomials, arising in combinatorics and in computer science. It is also known in statistical physics as the partition function of the hard-core model. This paper studies the computational complexity of evaluating the polynomial approximately when a parameter, called the *activity*, is complex. For properties of this polynomial in the complex plane, including connections to the Lovász Local Lemma, see the work of Scott and Sokal [11]. For $\lambda \in \mathbb{C}$ and a graph G the polynomial is defined as $Z_G(\lambda) := \sum_I \lambda^{|I|}$, where the sum ranges over all independent sets of G . We will be interested in the problem of approximating $Z_G(\lambda)$ when the maximum degree of G is bounded.

When λ is real, the complexity picture is well-understood. For $\Delta \geq 3$, let \mathcal{G}_Δ be the set of graphs with maximum degree at most Δ . The complexity of approximating $Z_G(\lambda)$ for $G \in \mathcal{G}_\Delta$ is captured by two real-valued thresholds λ^* and λ_c which depend on Δ and satisfy $0 < \lambda^* < \lambda_c$. To be precise, $\lambda^* = \frac{(\Delta-1)^{\Delta-1}}{\Delta^\Delta}$ and $\lambda_c = \frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^\Delta}$. The known results are as follows.

1. If λ is in the interval $-\lambda^* < \lambda < \lambda_c$, there is an FPTAS for approximating $Z_G(\lambda)$ on graphs $G \in \mathcal{G}_\Delta$. For $0 < \lambda < \lambda_c$, this follows from the work of Weitz [14], while for $-\lambda^* < \lambda < 0$ it follows from the works of Harvey, Srivastava, and Vondrák [6] and Patel and Regts [9].
2. If $\lambda < -\lambda^*$ or $\lambda > \lambda_c$, it is NP-hard to approximate $|Z_G(\lambda)|$ on graphs $G \in \mathcal{G}_\Delta$, even within an exponential factor. For $\lambda > \lambda_c$, this follows from the work of Sly and Sun [12], while for $\lambda < -\lambda^*$ it follows from the work of Galanis, Goldberg, and Štefankovič [4].

The key to establishing this complexity characterisation was the following interpretation of the thresholds λ^* and λ_c . Given a Δ -regular tree T of height h with root ρ , let p_h denote the “occupation ratio” of the tree, which is given by $p_h = \frac{\sum_{I: \rho \in I} \lambda^{|I|}}{Z_T(\lambda)}$, where the sum ranges over the independent sets of T that include the root ρ . It turns out that the occupation ratio p_h converges to a limit as $h \rightarrow \infty$ if and only if the activity λ lies within the interval $[-\lambda^*, \lambda_c]$, so the complexity of approximating $Z_G(\lambda)$ for $G \in \mathcal{G}_\Delta$ depends on whether this limit converges.

Understanding the complexity picture in the case where $\lambda \in \mathbb{C}$ is more challenging. If λ is a complex number with norm at most λ^* then there is an FPTAS for approximating $Z_G(\lambda)$ on graphs $G \in \mathcal{G}_\Delta$. This is due to Harvey, Srivastava and Vondrák and to Patel and Regts [6, 9]. More recently, Peters and Regts [10] showed the existence of an FPTAS when λ is in a small strip surrounding the real interval $[0, \lambda_c)$. However, neither of these results is believed to fully capture the truth about when approximation is possible. Motivated by the real case, Peters and Regts [10] identified the values of λ for which the occupation ratio of the Δ -regular tree converges (for $\Delta \geq 3$). These values carve a cardioid-shaped region Λ_Δ in the complex plane, whose boundary includes the critical points $-\lambda^*$ and λ_c . The definition of Λ_Δ is as follows (see Figure 1)¹:

$$\Lambda_\Delta = \left\{ \lambda \in \mathbb{C} \mid \exists z \in \mathbb{C} : |z| \leq 1/(\Delta-1), \lambda = \frac{z}{(1-z)^\Delta} \right\}. \quad (1)$$

Peters and Regts showed that, for every λ in the (strict) interior of Λ_Δ , the occupation ratio of the Δ -regular tree converges, and asked whether the region Λ_Δ marks the true approximability threshold for general complex values λ .

Our main result shows that for every λ outside of the region Λ_Δ , the problem of approximating $Z_G(\lambda)$ on graphs $G \in \mathcal{G}_\Delta$ is indeed NP-hard, thus answering [10, Question 1]. In fact, when λ is outside of Λ_Δ and is not a positive real number, we establish the stronger result that approximating $Z_G(\lambda)$ is actually #P-hard.

¹Technically, the word “cardioid” refers to a curve which can be obtained by a point on the perimeter of a circle which is rolling around a fixed circle of the same radius. The region (1) does not formally correspond to a “cardioid” in this sense, but its shape closely resembles a heart for all values of $\Delta \geq 3$, which justifies the (slight) abuse of terminology.




Figure 1: The cardioid-shaped region Λ_Δ in the complex plane. We show that for all $\lambda \in \mathbb{C} \setminus (\Lambda_\Delta \cup \mathbb{R}_{\geq 0})$, approximating $Z_G(\lambda)$ is $\#P$ -hard. Previously, it was known that the problem is NP-hard on the real line in the intervals $\lambda < -\lambda^*$ and $\lambda > \lambda_c$. Note, we have that the thresholds $-\lambda^*, \lambda_c$ belong to Λ_Δ , by taking $z = \pm 1/(\Delta - 1)$ in (1).

We do this by showing that an approximation algorithm for $Z_G(\lambda)$ can be converted into a polynomial-time algorithm for exactly counting independent sets. Further, on the negative real axis, when $\lambda < -\lambda^*$, we show that it is $\#P$ -hard to even decide whether $Z_G(\lambda) > 0$, resolving in the affirmative a conjecture of Harvey, Srivastava, and Vondrák [6, Conjecture 5.1].

We need the following notation to formally state our results. Given a complex number $x \in \mathbb{C}$, we use $|x|$ to denote its norm and $\text{Arg}(x)$ to denote the principal value of its argument in the range $[0, 2\pi)$. We also define $\arg(x) = \{\text{Arg}(x) + 2\pi j \mid j \in \mathbb{Z}\}$. For $y, z \in \mathbb{C}$, we use $d(y, z)$ to denote the Ziv distance between them [15], namely $d(y, z) = \frac{|y-z|}{\max(|y|, |z|)}$. We denote by $\mathbb{C}_{\mathbb{Q}}$ the set of complex numbers whose real and imaginary parts are rational numbers.

We consider the problems of multiplicatively approximating the norm of $Z_G(\lambda)$, additively approximating the argument of $Z_G(\lambda)$, and approximating $Z_G(\lambda)$ by producing a complex number \hat{Z} such that the Ziv distance $d(\hat{Z}, Z_G(\lambda))$ is small. We start with the following problem, which captures the approximation of the norm of $Z_G(\lambda)$.

Name $\#BipHardCoreNorm(\lambda, \Delta, K)$.

Instance A bipartite graph G with maximum degree at most Δ .

Output If $|Z_G(\lambda)| = 0$ then the algorithm may output any rational number. Otherwise, it must output a rational number \hat{N} such that $\hat{N}/K \leq |Z_G(\lambda)| \leq K\hat{N}$.

Our first theorem shows that it is $\#P$ -hard to approximate $|Z_G(\lambda)|$ on bipartite graphs of maximum degree Δ within a constant factor.

Theorem 1. *Let $\Delta \geq 3$ and $\lambda \in \mathbb{C}_{\mathbb{Q}}$ be such that $\lambda \notin (\Lambda_\Delta \cup \mathbb{R}_{\geq 0})$. Then, $\#BipHardCoreNorm(\lambda, \Delta, 1.01)$ is $\#P$ -hard.*

Remark. The value “1.01” in the statement of Theorem 1 is not important. In fact, for any fixed $\epsilon > 0$ we can use the theorem, together with a standard powering argument, to show that it is $\#P$ -hard to approximate $|Z_G(\lambda)|$ within a factor of $2^{n^{1-\epsilon}}$.

The following problem captures the approximation of the argument of $Z_G(\lambda)$.

Name $\#BipHardCoreArg(\lambda, \Delta, \rho)$.

Instance A bipartite graph G with maximum degree at most Δ .

Output If $Z_G(\lambda) = 0$ then the algorithm may output any rational number. Otherwise, it must output a rational number \hat{A} such that, for some $a \in \arg(Z_G(\lambda))$, $|\hat{A} - a| \leq \rho$.

Our second theorem shows that it is $\#P$ -hard to approximate $\arg(Z_G(\lambda))$ on bipartite graphs of maximum degree Δ within an additive constant $\pi/3$.

Theorem 2. Let $\Delta \geq 3$ and $\lambda \in \mathbb{C}_{\mathbb{Q}}$ be such that $\lambda \notin (\Lambda_{\Delta} \cup \mathbb{R}_{\geq 0})$. Then, $\#BipHardCoreArg(\lambda, \Delta, \pi/3)$ is $\#P$ -hard.

Theorem 2 also has the following immediate corollary for the case in which λ is a negative real number, resolving in the affirmative [6, Conjecture 5.1].

Corollary 3. Let $\Delta \geq 3$ and $\lambda \in \mathbb{Q}$ be such that $\lambda < -\lambda^*$. Then, given as input a bipartite graph G with maximum degree Δ , it is $\#P$ -hard to decide whether $Z_G(\lambda) > 0$.

Theorems 1 and 2 show as a corollary that it is $\#P$ -hard to approximate $Z_G(\lambda)$ within small Ziv distance, based on the fact that (see [5, Lemma 2.1]) that $d(z', z) \leq \epsilon$ implies $|z'|/|z| \leq 1/(1 - \epsilon)$, see full version for details.

Name $\#BipComplexHardCore(\lambda, \Delta)$.

Instance A bipartite graph G with maximum degree at most Δ . A positive integer R , in unary.

Output If $Z_G(\lambda) = 0$ then the algorithm may output any complex number. Otherwise, it must output a complex number z such that $d(z, Z_G(\lambda)) \leq 1/R$.

Corollary 4. Let $\Delta \geq 3$ and $\lambda \in \mathbb{C}_{\mathbb{Q}}$ be such that $\lambda \notin (\Lambda_{\Delta} \cup \mathbb{R}_{\geq 0})$. Then, $\#BipComplexHardCore(\lambda, \Delta)$ is $\#P$ -hard.

Note that our $\#P$ -hardness results for $\lambda \in \mathbb{C}_{\mathbb{Q}} \setminus (\Lambda_{\Delta} \cup \mathbb{R}_{\geq 0})$ highlight a difference in complexity between this case and the case where λ is a rational satisfying $\lambda > \lambda_c$. If λ is a positive rational then $Z_G(\lambda)$ can be efficiently approximated in polynomial time using an NP oracle, via the bisection technique of Valiant and Vazirani [13]. Thus, in that case approximation is NP-easy, and is unlikely to be $\#P$ -hard.

1.1 Proof approach

To prove our inapproximability results, we construct graph gadgets which, when appended appropriately to a vertex, have the effect of altering the activity λ to any complex activity λ' that we wish, perhaps with some small error ϵ . In fact, it is essential for our $\#P$ -hardness results to be able to make the error ϵ exponentially small with respect to the number of the vertices in the graph (see the upcoming Lemma 6 for details).

Interestingly, our constructions are based on using tools from complex analysis for analysing the iteration of rational maps. We start with the observation that $(\Delta - 1)$ -ary trees of height h can be used to “implement” activities λ' which correspond to the iterates of the complex rational map $f : x \mapsto \frac{1}{1 + \lambda x^{\Delta-1}}$. Crucially, we show that when $\lambda \notin \Lambda_{\Delta}$, all of the fixpoints of f are repelling, i.e., applying the map f at any point close

to a fixpoint ω will push us away from the fixpoint. In the iteration of univariate complex rational maps, repelling fixpoints belong to the so-called Julia set of the map; a consequence of this is that iterating f in a neighbourhood U of a repelling fixpoint gives rise to a chaotic behaviour: after sufficiently many iterations, one ends up anywhere in the complex plane.

This sounds promising, but how can we get close to a *repelling* fixpoint of f in the first place? In fact, we need to be able to create arbitrary points in a neighbourhood U of a repelling fixpoint and iterating f will not get us anywhere close (since the fixpoint is repelling). The key is to use a Fibonacci-type construction which requires analysing a more intricate multivariate version of the map f . Surprisingly, we can show that the iterates of the multivariate version converge to the fixpoint ω of the univariate f with the smallest norm. Using convergence properties of the multivariate map around ω (and some extra work), we obtain a family of (univariate) contracting² maps Φ_1, \dots, Φ_t and a small neighbourhood U around ω such that $U \subseteq \cup_{i=1}^t \Phi_i(U)$. The final step is to show that “contracting maps that cover yield exponential precision”. To do this we first show that, starting from any point in U , we can apply (some sequence of) Φ_1, \dots, Φ_t at most $\text{poly}(n)$ times to implement any point in U with precision $\exp(-\Omega(n))$. We then show that by iteratively applying the univariate map f and carefully tracking the distortion introduced, we can eventually implement any point in the complex plane with exponentially small error.

2 Proof Outline

In this section, we give a more detailed outline of the proof of our results. We focus mainly on the case where $\lambda \in \mathbb{C}_{\mathbb{Q}} \setminus (\Lambda_{\Delta} \cup \mathbb{R})$. In Section 2.7 we describe suitable modifications that will give us the ingredients needed for negative real values $\lambda \in \mathbb{Q} \setminus \Lambda_{\Delta}$.

Let $\lambda \in \mathbb{C}$ and $G = (V, E)$ be an arbitrary graph. We denote by \mathcal{I}_G the set of independent sets of G . For a vertex $v \in V$, we will denote

$$Z_{G,v}^{\text{in}}(\lambda) := \sum_{I \in \mathcal{I}_G; v \in I} \lambda^{|I|}, \quad Z_{G,v}^{\text{out}}(\lambda) := \sum_{I \in \mathcal{I}_G; v \notin I} \lambda^{|I|}.$$

Thus, $Z_{G,v}^{\text{in}}(\lambda)$ is the contribution to the partition function $Z_G(\lambda)$ from those independent sets $I \in \mathcal{I}_G$ such that $v \in I$; similarly, $Z_{G,v}^{\text{out}}(\lambda)$ is the contribution to $Z_G(\lambda)$ from those $I \in \mathcal{I}_G$ such that $v \notin I$.

Definition 5. Fix a complex number λ that is not 0. Given λ , the graph G is said to implement the activity $\lambda' \in \mathbb{C}$ with accuracy $\epsilon > 0$ if there is a vertex v in G such that $Z_{G,v}^{\text{out}}(\lambda) \neq 0$ and

1. v has degree one in G , and

2. $\left| \frac{Z_{G,v}^{\text{in}}(\lambda)}{Z_{G,v}^{\text{out}}(\lambda)} - \lambda' \right| \leq \epsilon.$

We call v the terminal of G . If Item 2 holds with $\epsilon = 0$, then G is said to implement the activity λ' .

The key to obtaining our #P-hardness results is to show that, given *any* target activity $\lambda' \in \mathbb{C}$, we can construct in polynomial time a bipartite graph G that implements λ' with exponentially small accuracy, as a function of the size of λ' . More precisely, we use $\text{size}(\lambda', \epsilon)$ to denote the number of bits needed to represent the complex number $\lambda' \in \mathbb{C}_{\mathbb{Q}}$ and the rational ϵ . The implementation that we need is captured by the following lemma.

²Let $\Phi : \mathbb{C} \rightarrow \mathbb{C}$ be a complex map. We say that Φ is contracting on a set $S \subseteq \mathbb{C}$ if there exists a real number $M < 1$ such that for all $x, y \in S$ it holds that $|\Phi(x) - \Phi(y)| \leq M|x - y|$.

Lemma 6. Let $\Delta \geq 3$ and $\lambda \in \mathbb{C}_{\mathbb{Q}}$ be such that $\lambda \notin \Lambda_{\Delta} \cup \mathbb{R}$.

There is an algorithm which, on input $\lambda' \in \mathbb{C}_{\mathbb{Q}}$ and rational $\epsilon > 0$, outputs in $\text{poly}(\text{size}(\lambda', \epsilon))$ time a bipartite graph G of maximum degree at most Δ with terminal v that implements λ' with accuracy ϵ . Moreover, the algorithm outputs the values $Z_{G,v}^{\text{in}}(\lambda)$, $Z_{G,v}^{\text{out}}(\lambda)$.

Lemma 6 is extremely helpful in our reductions since it enables us to construct other gadgets very easily, e.g., equality gadgets that reduce the degree of a graph and gadgets that can turn it into a bipartite graph. The proofs of Theorems 1 and 2 show how to use these gadget to obtain #P hardness. In this proof outline, we focus on the most difficult part which is the proof of Lemma 6.

To prove Lemma 6, we will make use of the following multivariate map:

$$(x_1, \dots, x_d) \mapsto \frac{1}{1 + \lambda x_1 \cdots x_d}, \text{ where } d := \Delta - 1. \quad (2)$$

If, starting from 1, there is a sequence of operations (2) which ends with the value x , for the purposes of this outline, we will loosely say that “we can generate the value x ” (the notion is formally defined in Definition 38 of the appendix). There is a simple correspondence between the values that we can generate and the activities that we can implement: in Lemma 39 of the appendix, we show that if we can generate a value x , we can also implement the activity λx using a tree of maximum degree Δ .³

To get some insight about the map (2), the first natural step is to look at the univariate case $x_1 = \dots = x_d = x$, where the map (2) simplifies into

$$f : x \mapsto \frac{1}{1 + \lambda x^d}.$$

Even analysing the iterates of this map is a surprisingly intricate task; fortunately there is a rich theory concerning the iteration of complex rational maps which we can use (though much less is known in the multivariate setting!). In the next section, we review the basic ingredients of the theory that we need, see [1, 7] for detailed accounts on the subject.

2.1 Iteration of complex rational maps

We will use $\widehat{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$ to denote the Riemann sphere (complex numbers with infinity). To handle ∞ , it will be convenient to consider the *chordal* metric $d(\cdot, \cdot)$ on the Riemann sphere $\widehat{\mathbb{C}}$, which is given for $z, w \in \mathbb{C}$ by

$$d(z, w) = \frac{2|z - w|}{(1 + |z|^2)^{1/2}(1 + |w|^2)^{1/2}}, \text{ and } d(z, \infty) = \lim_{w \rightarrow \infty} d(z, w) = \frac{2}{(1 + |z|^2)^{1/2}}.$$

Note that $d(z, w)$ is bounded by an absolute constant for all $z, w \in \widehat{\mathbb{C}}$.

Let $f : \widehat{\mathbb{C}} \rightarrow \widehat{\mathbb{C}}$ be a complex rational map, i.e., $f(z) = P(z)/Q(z)$ for some coprime polynomials P, Q . We define $f(\infty)$ as the limit of $f(z)$ when $z \rightarrow \infty$. The *degree* of f is the maximum of the degrees of P, Q . A point $p \in \mathbb{C}$ is called a *pole* of f if $Q(p) = 0$; when $p = \infty$, p is a pole of f if 0 is a pole of $f(1/z)$.

Suppose that $z^* \in \mathbb{C}$ is a fixpoint of f , i.e., $f(z^*) = z^*$. The multiplier of f at z^* is given by $q = f'(z^*)$. If $z^* = \infty$, the multiplier of f at z^* is given by $1/f'(\infty)$. Depending on the value of $|q|$, the fixpoint z^* is classified as follows: (i) *attracting* if $|q| < 1$, (ii) *repelling* if $|q| > 1$, and (iii) *neutral* if $|q| = 1$.

³Note the extra factor of λ when we pass to the implementation setting which is to ensure the degree requirement in Item 1 of Definition 5; while the reader should not bother at this stage with this technical detail, the statements of our lemmas are usually about implementing activities and therefore have this extra factor λ incorporated.

For a non-negative integer $n \geq 0$, we will denote by f^n the n -fold iterate of f (for $n = 0$, we let f^0 be the identity map). Given $z_0 \in \widehat{\mathbb{C}}$, the sequence of points $\{z_n\}$ defined by $z_n = f(z_{n-1}) = f^n(z_0)$ is called the *orbit* of z_0 .

Given a rational map $f : \widehat{\mathbb{C}} \rightarrow \widehat{\mathbb{C}}$, we will be interested in the sensitivity of an orbit under small perturbations of the starting point. A point z_0 belongs to the *Fatou* set if, for every $\epsilon > 0$ there exists $\delta > 0$ such that, for any point z' with $d(z', z_0) \leq \delta$, it holds that $d(f^n(z'), f^n(z_0)) \leq \epsilon$ for all positive integer n (in other words, z_0 belongs to the Fatou set if the family of maps $\{f^n\}_{n \geq 1}$ is equicontinuous at z_0 under the chordal metric). A point z_0 belongs to the *Julia* set if z_0 does not belong to the Fatou set (i.e., the Julia set is the complement of the Fatou set).

Lemma 7 (e.g., [7, Lemma 4.6]). *Every repelling fixpoint belongs to the Julia set.*

For $z \in \widehat{\mathbb{C}}$, the *grand orbit* $[z]$ is the set of points z' whose orbit intersects the orbit of z , i.e., for every $z' \in [z]$, there exist integers $m, n \geq 0$ such that $f^m(z) = f^n(z')$. The exceptional set of the map f is the set of points z whose grand orbit $[z]$ is finite. It turns out that the exceptional set of a rational map f can have at most two points and, in our applications, it will in fact be empty (see Lemma 23 of the appendix for details).

For $z_0 \in \mathbb{C}$ and $r > 0$, we use $B(z_0, r)$ to denote the ball of radius r around z_0 . A set U is a neighbourhood of z_0 if U contains a ball $B(z_0, r)$ for some $r > 0$. We will use the following fact.

Theorem 8 (see, e.g., [7, Theorem 4.10]). *Let $f : \widehat{\mathbb{C}} \rightarrow \widehat{\mathbb{C}}$ be a complex rational map with exceptional set E_f . Let z_0 be a point in the Julia set and let U be an arbitrary neighbourhood of z_0 . Then, the union of the forward images of U , i.e., the set $\bigcup_{n \geq 0} f^n(U)$, contains $\widehat{\mathbb{C}} \setminus E_f$.*

Peters and Regts [10] used a version of Theorem 8 to conclude the existence of trees T and λ 's close to the boundary of Λ_Δ such that $Z_T(\lambda) = 0$. We will use Theorem 8 as a tool to get our #P-hardness results for any λ outside the cardioid Λ_Δ .

2.2 A characterisation of the cardioid

To use the tools from the previous section, we will need to analyse the fixpoints of the map $f(z) = \frac{1}{1 + \lambda z^{\Delta-1}}$. We denote by $\vartheta\Lambda_\Delta$ the following curve (which is actually the boundary of the region Λ_Δ defined in (1))⁴:

$$\vartheta\Lambda_\Delta = \left\{ \lambda \in \mathbb{C} \mid \exists z \in \mathbb{C} : |z| = 1/(\Delta - 1), \lambda = \frac{z}{(1 - z)^\Delta} \right\}. \quad (3)$$

The following lemma is proved in Section 4 of the appendix.

Lemma 9. *Let $\Delta \geq 3$ and consider the map $f(z) = \frac{1}{1 + \lambda z^{\Delta-1}}$ for $\lambda \in \mathbb{C}$. Then,*

1. *For all $\lambda \in \Lambda_\Delta \setminus \vartheta\Lambda_\Delta$, f has a unique attractive fixpoint; all other fixpoints are repelling.*
2. *For all $\lambda \in \vartheta\Lambda_\Delta$, f has a unique neutral fixpoint; all other fixpoints are repelling.*
3. *For all $\lambda \notin \Lambda_\Delta$, all of the fixpoints of f are repelling.*

⁴The fact that the curve $\vartheta\Lambda_\Delta$, as defined in (3), is the boundary of the region Λ_Δ (defined in (1)) is shown in Footnote 4 of the appendix.

2.3 Applying the theory

We are now in a position to discuss in detail how to apply the tools of Section 2.1 and the result of Section 2.2.

Let $\lambda \in \mathbb{C}_{\mathbb{Q}} \setminus (\Lambda_{\Delta} \cup \mathbb{R})$. By Lemma 9, all of the fixpoints of $f = \frac{1}{1+\lambda z^{\Delta-1}}$ are repelling. By Lemma 7, all of the repelling fixpoints belong to the Julia set of the map, and therefore, by applying Theorem 8, iteratively applying f to a neighbourhood U of a repelling fixpoint gives the entire complex plane. Therefore, if we want to generate an arbitrary complex value $\lambda' \in \mathbb{C}$, it suffices to be able to generate values in a neighbourhood U close to a repelling fixpoint of f . Of course, in our setting we will also need to do this efficiently, up to exponential precision. The following lemma is therefore the next important milestone. It formalises exactly what we need to show in order to be able to prove Lemma 6.

Lemma 10. *Let $\Delta \geq 3$ and $\lambda \in \mathbb{C}_{\mathbb{Q}} \setminus \mathbb{R}$, and set $d := \Delta - 1$. Let ω be the fixpoint of $f(x) = \frac{1}{1+\lambda x^d}$ with the smallest norm.⁵ There exists a rational $\rho > 0$ such that the following holds.*

There is a polynomial-time algorithm such that, on input $\lambda' \in B(\lambda\omega, \rho) \cap \mathbb{C}_{\mathbb{Q}}$ and rational $\epsilon > 0$, outputs a bipartite graph G of maximum degree at most Δ with terminal v that implements λ' with accuracy ϵ . Moreover, the algorithm outputs the values $Z_{G,v}^{\text{in}}(\lambda)$, $Z_{G,v}^{\text{out}}(\lambda)$.

To briefly explain why Lemma 6 follows from Lemma 10, we first show how to use Lemma 10 to implement activities λx^* where x^* is close to a pole p of f (i.e., a point p which satisfies $1 + \lambda p^d = 0$). For some $r > 0$, let U be the ball $B(\omega, r)$ of radius r around ω . Using Theorem 8, we find the first integer value of $N > 0$ such that a pole p^* belongs to $f^N(U)$; in fact, we can choose r (see Lemma 29 of the appendix) so that there exists a radius $r^* > 0$ such that $B(p^*, r^*) \subseteq f^N(U)$. The idea of “waiting till we hit the pole of f ” is that, up to this point, the iterates of f satisfy Lipchitz inequalities, i.e., it can be shown that there exists a real number $L > 0$ such that $|f^N(x_1) - f^N(x_2)| \leq L|x_1 - x_2|$ for all $x_1, x_2 \in U$. Therefore, for any desired target $x^* \in B(p^*, r^*)$ we can find $w^* \in U$ such that $f^N(w^*) = x^*$ and implement λw^* using Lemma 10 with accuracy $\epsilon > 0$; due to the Lipchitz inequality, this yields an implementation of λx^* with accuracy at most $\lambda L\epsilon$, i.e., just a constant factor distortion. Once we are able to create specified activities close to λp^* where p^* is a pole of f , it is then possible to first implement activities λ' with large norm by plugging appropriate x close to p^* in $f(x) = \frac{1}{1+\lambda x^d}$; in turn, this can be then used to implement activities λ' with small norm and finally, λ' with moderate value of $|\lambda'|$ as well. See the proof of Lemma 6 in Section 5.3 of the appendix for more details.

2.4 Chasing repelling fixpoints

In this section, we focus on the proof of Lemma 10, whose proof (given in Section 7 of the appendix) requires us to delve into the analysis of the multivariate map

$$(x_1, \dots, x_d) \mapsto \frac{1}{1 + \lambda x_1 \cdots x_d}, \text{ where } d := \Delta - 1. \quad (2)$$

Recall, in the scope of proving Lemma 10, our goal is to generate points close to a repelling fixpoint of the map $f : x \mapsto \frac{1}{1+\lambda x^d}$. Since λ is outside the cardioid region Λ_{Δ} , the fixpoints of the map f are repelling and therefore we cannot get close to any of them by just iterating f . Can the multivariate map make it easier to get to a fixpoint of f ? The answer to the question is yes, as the following lemma asserts.

Lemma 11. *Let $\Delta \geq 3$ and $\lambda \in \mathbb{C} \setminus \mathbb{R}$, and set $d := \Delta - 1$. Let ω be the fixpoint of $f(x) = \frac{1}{1+\lambda x^d}$ with the smallest norm. For $k \geq 0$, let x_k be the sequence defined by*

$$x_0 = x_1 = \cdots = x_{d-1} = 1, \quad x_k = \frac{1}{1 + \lambda \prod_{i=1}^d x_{k-i}} \quad \text{for } k \geq d. \quad (4)$$

⁵Note, by Lemma 25 in the appendix, all the fixpoints of f have different norms for $\lambda \in \mathbb{C}_{\mathbb{Q}} \setminus \mathbb{R}$, so ω is well-defined.

Then, the sequence x_k is well-defined (i.e., the denominator of (4) is nonzero for all $k \geq d$) and converges to the fixpoint ω as $k \rightarrow \infty$. Moreover, there exist infinitely many k such that $x_k \neq \omega$.

Note, Lemma 9 guarantees that the fixpoint ω in Lemma 11 is repelling when $\lambda \in \mathbb{C} \setminus (\Lambda_\Delta \cup \mathbb{R})$, so Lemma 11 indeed succeeds in getting us close to a repelling fixpoint in this case. It is instructive at this point to note that the sequence in (4) corresponds to a Fibonacci-type tree construction T_0, \dots, T_k , where for $k \geq d$ tree T_k consists of a root r with subtrees T_{k-d}, \dots, T_{k-1} rooted at the children of r . The trees T_{k-d}, \dots, T_{k-1} generate the values x_{k-d}, \dots, x_{k-1} , respectively, and the tree T_k generates the value x_k .

A few remarks about the proof of Lemma 11 are in order. Analysing the behaviour of multivariate recurrences such as the one in (4) is typically an extremely complicated task and the theory for understanding such recurrences appears to be still under development. Fortunately, the recurrence (4) can be understood in a surprisingly simple way by using the linear recurrence R_k defined by $R_0 = \dots = R_d = 1$ and $R_{k+1} = R_k + \lambda R_{k-d}$ for all $k \geq d$, and observing that $x_k = R_k / R_{k+1}$ for all k . By interpreting R_k as the independent set polynomial of a claw-free graph evaluated at $\lambda \in \mathbb{C} \setminus \mathbb{R}$, we obtain using a result of Chudnovsky and Seymour [3] that $R_k \neq 0$. The detailed proof of Lemma 11 can be found in Section 7.1 of the appendix.

2.5 Contracting maps that cover yield exponential precision

Lemma 11 resolves the intriguing task of getting close to a repelling fixpoint ω of the univariate map when $\lambda \in \mathbb{C} \setminus (\Lambda_\Delta \cup \mathbb{R})$. But in the context of Lemma 10 we need to accomplish far more: we need to be able to generate any point which is in some (small) ball U around the fixpoint ω , with exponentially small error ϵ .

To do this, we will focus on a small ball U around ω , i.e., $U = B(\omega, \delta)$ for some sufficiently small $\delta > 0$, and we will examine how the multivariate map (2) behaves when $x_1, \dots, x_d \in U$. In particular, we show in Lemma 46 of the appendix that for any choice of $x_1, \dots, x_d \in U$ it holds that

$$\frac{1}{1 + \lambda x_1 \dots x_d} \approx \omega + z \left((x_1 - \omega) + \dots + (x_d - \omega) \right), \text{ where } z \text{ satisfies } 0 < |z| < 1 \text{ and } z \in \mathbb{C} \setminus \mathbb{R}. \quad (5)$$

The important observation is that once we fix $x_1, \dots, x_{d-1} \in B(\omega, \delta)$, the resulting map Φ is contracting with respect to the remaining argument x_d (in the vicinity of ω) — see Lemma 48 of the appendix for a more detailed treatment of this contraction.

The observation that Φ is contracting will form the basis of our approach to iteratively reduce the accuracy with which we need to generate points (by going backwards): if we need to generate a desired $x \in U$ with error at most ϵ it suffices to be able to generate $\Phi^{-1}(x)$ with error at most $\epsilon/|z| > \epsilon$, i.e., to generate x with good accuracy, we only need to do the easier task of generating the point $\Phi^{-1}(x)$ with less restrictive accuracy. The only trouble is that, if we use a single map Φ , after a few iterations of the process the preimage $\Phi^{-1}(x)$ will eventually escape U . To address this, note that in the construction of the map Φ above, we had the freedom to choose arbitrary $x_1, \dots, x_{d-1} \in B(\omega, \delta)$. We will make use of this freedom and, in particular, we will use a family of contracting maps Φ_1, \dots, Φ_t (for some large constant t) instead of a single map Φ ; the large number of maps will allow us to guarantee that for all $x \in U$, at least one of the preimages $\Phi_1^{-1}(x), \dots, \Phi_t^{-1}(x)$ belongs in U , i.e., that the images $\Phi_1(U), \dots, \Phi_t(U)$ cover U . We will discuss in Section 2.6 how to obtain the maps Φ_1, \dots, Φ_t , but first let us formalise the above into the following lemma, which is the basis of our technique for making the error exponentially small.

Lemma 12. *Let $z_0 \in \mathbb{C}_0$, $r > 0$ be a rational and U be the ball $B(z_0, r)$. Further, suppose that $\lambda'_1, \dots, \lambda'_t \in \mathbb{C}_0$ are such that the complex maps $\Phi_i : z \mapsto \frac{1}{1 + \lambda'_i z}$ with $i \in [t]$ satisfy the following:*

1. *for each $i \in [t]$, Φ_i is contracting on the ball U ,*
2. $U \subseteq \bigcup_{i=1}^t \Phi_i(U)$.

There is an algorithm which, on input (i) a starting point $x_0 \in U \cap \mathbb{C}_{\mathbb{Q}}$, (ii) a target $x \in U \cap \mathbb{C}_{\mathbb{Q}}$, and (iii) a rational $\epsilon > 0$, outputs in $\text{poly}(\text{size}(x_0, x, \epsilon))$ time a number $\hat{x} \in U \cap \mathbb{C}_{\mathbb{Q}}$ and a sequence $i_1, i_2, \dots, i_k \in [t]$ such that

$$\hat{x} = \Phi_{i_k}(\Phi_{i_{k-1}}(\dots \Phi_{i_1}(x_0) \dots)) \text{ and } |\hat{x} - x| \leq \epsilon.$$

The proof of Lemma 12 can be carried out along the lines we sketched above, see Section 5.1 of the appendix for details. In that section, we also pair Lemma 12 with a path construction which, given the sequence of indices i_1, \dots, i_k , returns a path of length k that implements $\lambda \hat{x}$ (cf. footnote 3 for the extra factor of λ), see Lemma 27 of the appendix for details.

2.6 Constructing the maps

We next turn to the last missing piece, which is to create the maps Φ_1, \dots, Φ_t which satisfy the hypotheses of Lemma 12 in a ball $U = B(\omega, \delta)$ around the fixpoint ω for some small radius $\delta > 0$ (note, we are free to make δ as small as we wish). The following notions of “covering” and “density” will be relevant for this section.

Definition 13. Let $U \subseteq \mathbb{C}$. A set $F \subseteq U$ is called an ϵ -covering of U if for every $x \in U$ there exists $y \in F$ such that $|x - y| \leq \epsilon$. A set $F \subseteq U$ is called dense in U if F is an ϵ -covering of U for every $\epsilon > 0$.

We have already seen in Section 2.5 that, for arbitrary $x_1, \dots, x_d \in U$, we have

$$\frac{1}{1 + \lambda x_1 \dots x_d} \approx \omega + z \left((x_1 - \omega) + \dots + (x_d - \omega) \right), \text{ where } z \text{ satisfies } 0 < |z| < 1 \text{ and } z \in \mathbb{C} \setminus \mathbb{R}. \quad (5)$$

We also discussed that, if we fix arbitrary $x_1, \dots, x_{d-1} \in U$, the resulting map $\Phi(x) = \frac{1}{1 + (\lambda x_1 \dots x_{d-1})x}$ is contracting in U for all sufficiently small $\delta > 0$, and therefore we can easily take care of the contraction properties that we need (in the context of Lemma 12). The more difficult part is to control the preimage of the map Φ . We show in Lemma 47 of the appendix that for $x, x_1, \dots, x_{d-1} \in U$, it holds that

$$\Phi^{-1}(x) = \frac{1}{\lambda x_1 \dots x_{d-1}} \left(\frac{1}{x} - 1 \right) \approx \omega + \left(\frac{x - \omega}{z} - \sum_{j=1}^{d-1} (x_j - \omega) \right).$$

Therefore to ensure that $\Phi^{-1}(x)$ belongs to $U = B(\omega, \delta)$ we need to ensure that x_1, \dots, x_{d-1} are such that

$$\left| \frac{x - \omega}{z} - \sum_{j=1}^{d-1} (x_j - \omega) \right| < \delta/2. \quad (6)$$

Note that by Lemma 11 we can generate points arbitrarily close to ω and hence we can make each of $x_2 - \omega, \dots, x_{d-1} - \omega$ so small that they are effectively negligible in (6); then, to be able to satisfy (6), we need to be able to choose x_1 so that $|(x - \omega)/z - (x_1 - \omega)|$ is small, say less than $\delta/4$. Since $|(x - \omega)/z| \leq \delta/|z|$, the key will therefore be to produce a $(\delta/4)$ -covering of the slightly enlarged ball $B(\omega, \delta/|z|)$. Then, we can take x_1 to be one of the points in the $(\delta/4)$ -covering.

We will in fact show the following slightly more general lemma, which guarantees that we can indeed generate the required points around ω for any desired precision $\epsilon > 0$ provided that we choose δ small enough (and can therefore implement activities around $\lambda\omega$). Note that the lemma can be viewed as a “relaxed” version of Lemma 10 with much weaker guarantees.

Lemma 14. Let $\Delta \geq 3$ and $\lambda \in \mathbb{C}_{\mathbb{Q}} \setminus \mathbb{R}$, and set $d := \Delta - 1$. Let ω be the fixpoint of $f(x) = \frac{1}{1 + \lambda x^d}$ with the smallest norm. For any $\epsilon, \kappa > 0$ there exists a radius $\rho \in (0, \kappa)$ such that the following holds. For every $\lambda' \in B(\lambda\omega, \rho)$, there exists a tree G of maximum degree at most Δ that implements λ' with accuracy $\rho\epsilon$.

But how can we “populate” the vicinity of ω , i.e., generate a covering of a ball $U = B(\omega, \delta)$? Lemma 11 only gives us that we can generate points arbitrarily close to ω . The key once again is to use the multivariate map around ω and, in particular, the perturbation estimate in the r.h.s of (5). To focus on the displacement from ω , we will use the transformation $a_i = x_i - \omega$ so that (5) translates into the following operation

$$(a_1, \dots, a_d) \mapsto z(a_1 + \dots + a_d),$$

i.e., if we have generated points which are displaced by a_1, \dots, a_d from ω , we can also generate a point which is roughly displaced by $z(a_1 + \dots + a_d)$ from ω ; we will only need to apply the operation a finite number of times, so the error coming from (5) will not matter critically and can be ignored in the following. We show in Lemma 45 of the appendix that, using a sequence of such operations, we can generate points of the form $\omega + z^{N(p)}p(z)$ where p is an arbitrary polynomial with non-negative integer coefficients and $N(p)$ is a positive integer which is determined by the number of operations we used to create p . We further show in Lemma 42 of the appendix, that for all $z \in \mathbb{C} \setminus \mathbb{R}$ with $|z| < 1$, the values $p(z)$, as p ranges over all polynomials with non-negative integer coefficients, form a dense set of \mathbb{C} . Therefore, to obtain Lemma 14, we can choose an ϵ -covering F of the unit disc using a finite set of values $p(z)$ and set $\delta = z^N$ where $N = \max_{p \in F} N(p)$; then, we can generate the points $\omega + \delta p(z)$ for every $p \in F$, which form an $(\epsilon\delta)$ -covering of the ball $U = B(\omega, \delta)$, yielding Lemma 14. The full proof is in Section 7.2 of the appendix.

2.7 Fitting the pieces together and proof for the real case

We briefly summarise the proof of Lemma 10. First, we get points close to a repelling fixpoint by showing Lemma 11 (discussed in Section 2.4 and proved in Section 7.1 of the appendix). Then, we bootstrap this into a moderately dense set of points around the fixpoint, yielding Lemma 14 (discussed in Section 2.6 and proved in Section 7.2). Further, we bootstrap this into exponential precision around the fixpoint, yielding Lemma 10 (discussed in Section 2.5 and proved in Section 7.3). Finally, we propagate this exponential precision to the whole complex plane, therefore yielding Lemma 10 (discussed in Section 2.3 and proved in Section 5.3).

Finally, we mention the modifications needed for the real case when $\lambda < -\lambda^*$. The following lemma is the analogue of Lemma 6 and allows us to implement real activities with exponential precision.

Lemma 15. *Let $\Delta \geq 3$ and $\lambda \in \mathbb{Q}$ be such that $\lambda < -\lambda^*$.*

There is an algorithm which, on input $\lambda', \epsilon \in \mathbb{Q}$ with $\epsilon > 0$, outputs in $\text{poly}(\text{size}(\lambda'), \epsilon)$ time a bipartite graph G of maximum degree at most Δ with terminal v that implements λ' with accuracy ϵ . Moreover, the algorithm outputs the values $Z_{G,v}^{\text{in}}(\lambda), Z_{G,v}^{\text{out}}(\lambda)$.

As in the complex case, we will need a moderately dense set of activities to get started, i.e., an analogue of Lemma 14; here, our job is somewhat simplified (relative to the case where $\lambda \in \mathbb{C} \setminus \mathbb{R}$) since we can use the following result of [4].

Lemma 16 ([4, Lemma 4]). *Let $\Delta \geq 3$ and $\lambda < -\lambda^*$. Then, for every $\lambda' \in \mathbb{R}$, for every $\epsilon > 0$, there exists a bipartite graph G of maximum degree at most Δ that implements λ' with accuracy ϵ .*

Note that Lemma 16 does not control the size of the graph G with respect to the accuracy ϵ , so it does not suffice to prove Lemma 15 on its own. In order to do this, we use the “contracting maps that cover” technique to get the exponential precision, i.e., the analogue of Lemma 12 restricted to the reals (see Lemma 26 of the appendix). The proof of Lemma 15 is completed in Section 5.2 of the appendix.

Once the proofs of Lemmas 6 and 15 are in place, we give the proofs of our #P-hardness results in Section 6 of the appendix.

References

- [1] A. F. Beardon. *Iteration of rational functions: complex analytic dynamical systems*. Graduate texts in mathematics. Springer-Verlag, 1991.
- [2] Y. Bugeaud. *Approximation by Algebraic Numbers*. Cambridge Tracts in Mathematics. Cambridge University Press, 2004.
- [3] Maria Chudnovsky and Paul D. Seymour. The roots of the independence polynomial of a clawfree graph. *J. Comb. Theory, Ser. B*, 97(3):350–357, 2007.
- [4] Andreas Galanis, Leslie Ann Goldberg, and Daniel Štefankovič. Inapproximability of the independent set polynomial below the Shearer threshold. In *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland*, pages 28:1–28:13, 2017.
- [5] Leslie Ann Goldberg and Heng Guo. The complexity of approximating complex-valued ising and tutte partition functions with applications to quantum simulation. *CoRR*, abs/1409.5627, 2014.
- [6] Nicholas J. A. Harvey, Piyush Srivastava, and Jan Vondrák. Computing the independence polynomial in Shearer’s region for the LLL. *CoRR*, abs/1608.02282, 2016.
- [7] J. W. Milnor. *Dynamics in One Complex Variable*. Number v. 10 in Academic Search Complete. Princeton University Press, 2006.
- [8] C. Andrew Neff and John H. Reif. An efficient algorithm for the complex roots problem. *Journal of Complexity*, 12(2):81 – 115, 1996.
- [9] Viresh Patel and Guus Regts. Deterministic polynomial-time approximation algorithms for partition functions and graph polynomials. *Electronic Notes in Discrete Mathematics*, 61(Supplement C):971 – 977, 2017. The European Conference on Combinatorics, Graph Theory and Applications (EURO-COMB’17).
- [10] Han Peters and Guus Regts. On a conjecture of Sokal concerning roots of the independence polynomial. *CoRR*, abs/1701.08049, 2017.
- [11] Alexander D. Scott and Alan D. Sokal. The repulsive lattice gas, the independent-set polynomial, and the Lovász local lemma. *J. Stat. Phys.*, 118(5-6):1151–1261, 2005.
- [12] Allan Sly and Nike Sun. Counting in two-spin models on d-regular graphs. *Ann. Probab.*, 42(6):2383–2416, 11 2014.
- [13] Leslie G. Valiant and Vijay V. Vazirani. NP is as easy as detecting unique solutions. *Theor. Comput. Sci.*, 47(3):85–93, 1986.
- [14] Dror Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the Thirty-eighth Annual ACM Symposium on Theory of Computing, STOC ’06*, pages 140–149, New York, NY, USA, 2006. ACM.
- [15] Abraham Ziv. Relative distance—an error measure in round-off error analysis. *Math. Comp.*, 39(160):563–569, 1982.