Scaling limits in algebra, geometry, and probability

Arzhakova, E.

Chapter 3

Decimation limits of principal algebraic $\mathbb{Z}^d$-actions[1]

Abstract

Let f be a Laurent polynomial in d commuting variables with integer coefficients. Associated to f is the principal algebraic $\mathbb{Z}^d$-action α_f on a compact subgroup X_f of $\mathbb{T}^{\mathbb{Z}^d}$ determined by f. Let $N \geqslant 1$ and restrict points in X_f to coordinates in $N\mathbb{Z}^d$. The resulting algebraic $N\mathbb{Z}^d$-action is again principal, and is associated to a polynomial g_N whose support grows with N and whose coefficients grow exponentially with N. We prove that by suitably renormalizing these decimations we can identify a limiting behavior given by a continuous concave function on the Newton polytope of f, and show that this decimation limit is the negative of the Legendre dual of the Ronkin function of f. In certain cases with two variables, the decimation limit coincides with the surface tension of random surfaces related to dimer models, but the statistical physics methods used to prove this are quite different and depend on special properties of the polynomial.

3.1 Introduction

Let $d \geqslant 1$ and $f \in \mathbb{Z}[x_1^{\pm 1}, \ldots, x_d^{\pm 1}]$ be a Laurent polynomial with integer coefficients in d commuting variables. We write $f(x_1, \ldots, x_d) = f(\mathbf{x}) =$

[1]This chapter is based on: E. Arzhakova, D. Lind, K. Schmidt, E. Verbitskiy, Decimation limits of principal algebraic $\mathbb{Z}^d$-actions, arXiv:2104.04408

$\sum_{\mathbf{n}\in\mathbb{Z}^d} \widehat{f}(\mathbf{n})\mathbf{x}^{\mathbf{n}}$, where $\mathbf{x}^{\mathbf{n}} = x_1^{n_1}\dots x_d^{n_d}$ and $\widehat{f}(\mathbf{n}) \in \mathbb{Z}$ for all $\mathbf{n} \in \mathbb{Z}^d$ and is nonzero for only finitely many $\mathbf{n} \in \mathbb{Z}^d$.

Denote the additive torus $\mathbb{R}/\mathbb{Z}$ by $\mathbb{T}$. Use f to define a compact subgroup X_f of $\mathbb{T}^{\mathbb{Z}^d}$ by

$$X_f := \Big\{t \in \mathbb{T}^{\mathbb{Z}^d} : \sum_{\mathbf{n}\in\mathbb{Z}^d} \widehat{f}(\mathbf{n})t_{\mathbf{m}+\mathbf{n}} = 0 \quad \text{for all } \mathbf{m} \in \mathbb{Z}^d\Big\}. \tag{3.1}$$

By its definition this subgroup is invariant under the natural shift-action σ of $\mathbb{Z}^d$ on $\mathbb{T}^{\mathbb{Z}^d}$ defined by $\sigma^{\mathbf{n}}(t)_{\mathbf{m}} = t_{\mathbf{m}-\mathbf{n}}$. Hence the restriction α_f of σ to X_f gives an action of $\mathbb{Z}^d$ by automorphisms of the compact abelian group X_f. We call (X_f, α_f) the *principal algebraic $\mathbb{Z}^d$-action defined by f*.

Such $\mathbb{Z}^d$-actions serve as a rich class of examples and have been studied intensively. An observation of Halmos [46] shows that α_f automatically preserves Haar measure μ_f on X_f. It is known that the topological entropy of α_f coincides with its measure-theoretic entropy with respect to μ_f. For nonzero f this common value was computed in [64] to be the logarithmic Mahler measure of f, defined as

$$m(f) := \int_0^1 \cdots \int_0^1 \log|f(e^{2\pi i s_1}, \dots, e^{2\pi i s_d})|\, ds_1 \dots ds_d \tag{3.2}$$

(when $f = 0$ the entropy is infinite).

It will be convenient to identify the Laurent polynomial ring $\mathbb{Z}[x_1^{\pm1}, \dots, x_d^{\pm1}]$ with the integral group ring $\mathbb{Z}[\mathbb{Z}^d]$, where the monomial $\mathbf{x}^{\mathbf{n}}$ corresponds to $\mathbf{n} \in \mathbb{Z}^d$. Thus $f \in \mathbb{Z}[x_1^{\pm1}, \dots, x_d^{\pm1}]$ is identified with its coefficient function $\widehat{f}\colon \mathbb{Z}^d \to \mathbb{Z}$. When emphasizing the behavior of coefficients we will always use the notation $\widehat{f}$.

Fix a principal algebraic $\mathbb{Z}^d$-action (X_f, α_f). Let $N \geqslant 1$ and $r_N\colon \mathbb{T}^{\mathbb{Z}^d} \to \mathbb{T}^{N\mathbb{Z}^d}$ be the map restricting the coordinates of a point to only those in the sublattice $N\mathbb{Z}^d$. We call the image $r_N(X_f)$ the *Nth decimation* of X_f, although this is considerably more brutal that the term's original meaning since only every Nth coordinate survives. Clearly $r_N(X_f)$ is again a compact abelian group, and it is invariant under the natural shift action of $N\mathbb{Z}^d$ on $\mathbb{T}^{N\mathbb{Z}^d}$.

Using commutative algebra applied to contracted ideals in integral extensions, we show in §3.6 that $r_N(X_f)$ is a principal algebraic $N\mathbb{Z}^d$-action with some defining polynomial $g_N \in \mathbb{Z}[N\mathbb{Z}^d]$. Typically the support of g_N grows with N and its coefficient function $\widehat{g}_N$ grows exponentially in N. Our goal in this paper is to prove that with suitable renormalizations the concave hulls

of the resulting functions converge uniformly on the Newton polytope of f to a continuous decimation limit D_f. Furthermore, D_f can be computed via Legendre duality using a well-studied object called the Ronkin function of f.

The analytical parts of our analysis apply to Laurent polynomials with complex coefficients. For such an $f \in \mathbb{C}[\mathbb{Z}^d]$ we define its Nth *decimation* f_N by

$$f_N(x_1, \ldots, x_d) := \prod_{k_1=0}^{N-1} \cdots \prod_{k_d=0}^{N-1} f(e^{2\pi i k_1/N} x_1, \ldots, e^{2\pi i k_d/N} x_d). \tag{3.3}$$

Since f_N is unchanged after multiplying each of its variables by an arbitrary Nth root of unity, it follows that it is a polynomial in the Nth powers of the x_i, i.e., that $f_N \in \mathbb{C}[N\mathbb{Z}^d]$. Decimations of polynomials have appeared in many contexts, including Purbhoo's approximations to shapes of complex amoebas [92], Boyd's proof that the Mahler measure of a polynomial is continuous in its coefficients [11], and dimer models in statistical physics [55].

For most $f \in \mathbb{Z}[\mathbb{Z}^d]$ the generator g_N of the Nth decimation of X_f coincides with f_N. But under special circumstances characterized in §3.6, involving the support of f and the Galois properties of the coefficients of the polynomials occurring in the factorization of f over the algebraic closure of the rationals, it can happen that g_N is a proper divisor of f_N. To give a simple example when $d = 1$, let $f(x) = x^2 - 2$. Then since f is already in $\mathbb{Z}[2\mathbb{Z}]$ we have that $g_2(x) = f(x)$, while $f_2(x) = f(x)f(-x) = f(x)^2$. Nevertheless even in these circumstances the renormalization behavior of the g_N can be determined from that of the f_N.

For $f \in \mathbb{C}[\mathbb{Z}^d]$ let $\operatorname{supp} f = \{\mathbf{n} \in \mathbb{Z}^d : \widehat{f}(\mathbf{n}) \neq 0\}$ denote its support. The *Newton polytope* $\mathcal{N}_f$ of f is the convex hull in $\mathbb{R}^d$ of $\operatorname{supp} f$. Since f_N is the product of N^d polynomials all of whose Newton polytopes are $\mathcal{N}_f$, it follows that $\mathcal{N}_{f_N} = N^d \mathcal{N}_f$.

The *Ronkin function* $R_f : \mathbb{R}^d \to \mathbb{R}$ of $0 \neq f \in \mathbb{C}[\mathbb{Z}^d]$ is defined by

$$R_f(u_1, \ldots, u_d) := \int_0^1 \cdots \int_0^1 \log |f(e^{u_1} e^{2\pi i s_1}, \ldots, e^{u_d} e^{2\pi i s_d})| \, ds_1 \ldots ds_d. \tag{3.4}$$

This is a convex function on $\mathbb{R}^d$, and therefore has a Legendre dual R_f^* defined by

$$R_f^*(\mathbf{r}) := \sup\{\mathbf{r} \cdot \mathbf{u} - R_f(\mathbf{u}) : \mathbf{u} \in \mathbb{R}^d\},$$

which turns out to be a convex function on $\mathcal{N}_f$ (and is ∞ off $\mathcal{N}_f$).

To describe rescaling of polynomials $g \in \mathbb{C}[\mathbb{Z}^d]$ it is convenient to extend the domain of $\widehat{g}$ from $\mathbb{Z}^d$ to $\mathbb{R}^d$ by declaring its value to be 0 off $\operatorname{supp} g$.

Let $\varphi\colon \mathbb{R}^d \to \mathbb{C}$. For any $a > 0$ define the *rescaling operator* E_a on φ by $(E_a\varphi)(\mathbf{r}) = \varphi(a\mathbf{r})$ for all $\mathbf{r} \in \mathbb{R}^d$. When dealing with concave functions it is often convenient to use the extended range $\underline{\mathbb{R}} = \mathbb{R} \cup \{-\infty\}$, with the usual algebraic rules for handling $-\infty$ and with the convention that $\log 0 = -\infty$. Then $\log|\varphi|\colon \mathbb{R}^d \to \underline{\mathbb{R}}$, and we define its *concave hull* $CH(\log|\varphi|)$ to be the infimum of all affine functions on $\mathbb{R}^d$ that dominate $\log|\varphi|$.

Let $f \in \mathbb{C}[\mathbb{Z}^d]$ and f_N be its Nth decimation. Define the Nth *logarithmic rescaling* $L_N f$ of f by

$$L_N f := E_{N^d}\Big(\frac{1}{N^d}\log|\widehat{f}_N|\Big).$$

Clearly $L_N f(\mathbf{r}) = -\infty$ if $\mathbf{r} \notin \mathcal{N}_f$, and is finite at every extreme point of $\mathcal{N}_f$ and at only finitely many other points in $\mathcal{N}_f$. The Nth *renormalized decimation* $D_N f$ of f is the concave hull $CH(L_N f)$ of $L_N f$. By our previous remark, $D_N f$ equals $-\infty$ off $\mathcal{N}_f$ and is finite at every point of $\mathcal{N}_f$.

With these preparations we can now state one of our main results.

Theorem 3.1.1. Let $0 \neq f \in \mathbb{C}[\mathbb{Z}^d]$. Then the Nth renormalized decimations $D_N f$ of f are concave polyhedral functions on the Newton polytope $\mathcal{N}_f$ of f that converge uniformly on $\mathcal{N}_f$ as $N \to \infty$ to a continuous concave decimation limit function D_f (and off $\mathcal{N}_f$ they are equal to $-\infty$). Furthermore $D_f = -R_f^*$, where R_f^* is the Legendre dual of the Ronkin function R_f of f.

The proof of this theorem uses two main ideas: Mahler's fundamental estimate [77] relating the largest coefficient of a polynomial to its Mahler measure and support, and a method used by Boyd [11], applied to decimations along powers of 2, to prove that for polynomials whose support is contained in a fixed finite subset of $\mathbb{Z}^d$ the Mahler measure is a continuous function of their coefficients.

If $f \in \mathbb{Z}[\mathbb{Z}^d]$ the decimation limit of f contains dynamical information about α_f.

Corollary 3.1.2. Let $0 \neq f \in \mathbb{C}[\mathbb{Z}^d]$. Then the maximum value of the decimation limit D_f on the Newton polytope $\mathcal{N}_f$ equals the logarithmic Mahler measure $m(f)$ of f defined in (3.2). In particular, if $f \in \mathbb{Z}[\mathbb{Z}^d]$ then this maximum value equals the entropy of the principal algebraic $\mathbb{Z}^d$-action α_f.

Duality allows us to compute the decimation limit of a product of two

polynomials. Suppose that $\varphi, \psi \colon \mathbb{R}^d \to \underline{\mathbb{R}}$ both have finite supremum. Define their *tropical convolution* $\varphi \circledast \psi$ by

$$(\varphi \circledast \psi)(\mathbf{r}) := \sup\{\varphi(\mathbf{s}) + \psi(\mathbf{r} - \mathbf{s}) : \mathbf{s} \in \mathbb{R}^d\}.$$

This is the tropical analogue of standard convolution, but using tropical (or max-plus) arithmetic in $\underline{\mathbb{R}}$.

Corollary 3.1.3. Let f and g be nonzero polynomials in $\mathbb{C}[\mathbb{Z}^d]$. Then $D_{fg} = D_f \circledast D_g$.

Thus decimation limits live in the tropics.

3.2 Examples

Here we give some examples to illustrate the phenomenon we are investigating. They use either one or two variables, and for these we denote the variables by x and y rather than x_1 and x_2. Let $\Omega_N = \{e^{2\pi ik/N} : 0 \leqslant k < N\}$ denote the group of Nth roots of unity.

Example 3.2.1. Let $d = 1$ and $f(x) = x^2 - x - 1 = (x - \lambda)(x - \mu)$, where $\lambda = (1 + \sqrt{5})/2$ and $\mu = (1 - \sqrt{5})/2$. Then

$$\begin{aligned} f_N(x) &= \prod_{\omega \in \Omega_N} f(\omega x) = \prod_{\omega \in \Omega_N} (\omega x - \lambda)(\omega x - \mu) \\ &= (x^N - \lambda^N)(x^N - \mu^N) = x^{2N} - (\lambda^n + \mu^N)x^N + (-1)^N. \end{aligned}$$

Hence

$$(L_N f)(r) = \begin{cases} 0 & \text{if } r = 0 \text{ or } 2, \\ \frac{1}{N} \log|\lambda^N + \mu^N| & \text{if } r = 1, \\ -\infty & \text{otherwise.} \end{cases}$$

Since $L_N f(1) \to \log \lambda$ as $N \to \infty$, the concave hulls $D_N f$ converge uniformly on $\mathcal{N}_f = [0, 2]$ to the decimation limit

$$D_f(r) = \begin{cases} r \log \lambda & \text{if } 0 \leqslant r \leqslant 1, \\ (2 - r) \log \lambda & \text{if } 1 \leqslant r \leqslant 2, \\ -\infty & \text{otherwise,} \end{cases}$$

which is shown in Figure 3.1(a).

To compute the Ronkin function R_f, recall Jensen's formula that for every $\xi \in \mathbb{C}$ we have that

$$\int_0^1 \log|e^{2\pi i s} - \xi|\, ds = \max\{0, \log|\xi|\} := \log^+|\xi|. \tag{3.5}$$

Thus

$$\begin{aligned} R_f(u) = \int_0^1 \log|f(e^u e^{2\pi i s})|\, ds &= \int_0^1 \log|e^u e^{2\pi i s} - \lambda|\, ds + \int_0^1 \log|e^u e^{2\pi i s} - \mu|\, ds \\ &= 2u + \log^+|e^{-u}\lambda| + \log^+|e^{-u}\mu|, \end{aligned}$$

whose polygonal graph is depicted in Figure 3.1(b). It is then easy to verify using the definition of Legendre transform that $D_f = -R_f^*$.

Finally, the decimation limits $D_{x-\lambda}$ and $D_{x-\mu}$ are computed similarly, and shown in Figures 3.1(c) and 3.1(d). It is easy to check using the definition of tropical convolution that $D_{x-\lambda} \circledast D_{x-\mu} = D_{(x-\lambda)(x-\mu)} = D_f$, in agreement with Corollary 3.1.3.

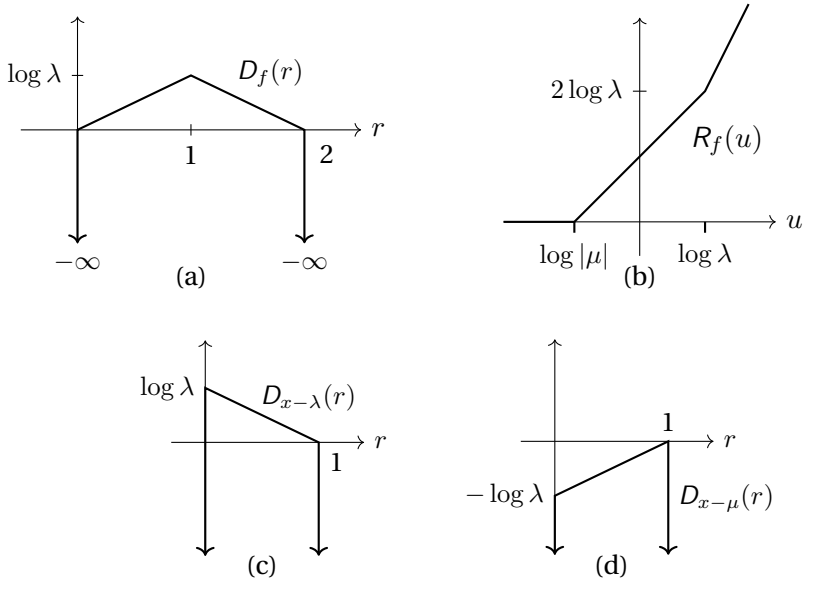

Figure 3.1: Graphs in Example 3.2.1

More generally, if $f(x) = \prod_{j=1}^m (x - \lambda_j)$ and $|\lambda_1| > |\lambda_2| > \cdots > |\lambda_m|$, then a computation similar to that in Example 3.2.1 shows that $(L_N f)(m) = 0$ and that $(L_N f)(k)$ converges to $\log|\lambda_1\lambda_2 \ldots \lambda_{m-k}|$ for $k = 0, 1, \ldots, m-1$, and this gives uniform convergence of $D_N f$ to D_f on $\mathcal{N}_f = [0, m]$. However, if some

roots of f have equal absolute value, then convergence is more delicate, or may even fail, as the next two examples show.

Example 3.2.2. Let $d = 1$ and $f(x) = x^4 - 4x^3 - 2x^2 - 4x + 1$, which is irreducible in $\mathbb{Z}[\mathbb{Z}]$. The roots of f are $\lambda = 1 + \sqrt{2} + \sqrt{2\sqrt{2}+2} \approx 4.611$, $\mu = 1 + \sqrt{2} - \sqrt{2\sqrt{2}+2} \approx 0.217$, and $1 - \sqrt{2} \pm i\sqrt{2\sqrt{2}-2} = e^{\pm 2\pi i\theta}$, where θ is irrational. Simple estimates show that $(L_N f)(k)$ converges for $k = 0, 1, 3, 4$ with limits $0, \log\lambda, \log\lambda, 0$, respectively. However, the dominant term controlling the behavior of $(L_N f)(2)$ is

$$\frac{1}{N}\log|2\lambda^N \cos(2\pi N\theta)|.$$

Since θ is irrational, the factor $\cos(2\pi N\theta)$ occasionally becomes very small, and so convergence is in question.

In fact, $(L_N f)(2)$ does converge, but the proof requires a deep result of Gelfond [37, Thm. III, p. 28] on the diophantine properties of algebraic numbers on the unit circle. According to this result, if ξ is an algebraic number (such as $e^{2\pi i\theta}$ above) such that $|\xi| = 1$ and ξ is not a root of unity, and if $\varepsilon > 0$, then $|\xi^n - 1| > e^{-n\varepsilon}$ for all but finitely many n. From this it is easy to deduce that $|e^{2\pi iN\theta} - i| > e^{-N\varepsilon}$ for almost every N, and hence that $(1/N)\log|\cos(2\pi N\theta)| \to 0$ as $N \to \infty$. This convergence is illustrated in Figure 3.2(a).

Both $(L_N f)(1)$ and $(L_N f)(3)$ converge to $\log\lambda$, and clearly it holds that $\limsup_{N\to\infty}(L_N f)(2) \leqslant \log\lambda$. Hence any lack of convergence of $(L_N f)(2)$ would not affect the limiting behavior of the concave hull $D_N f$, nor uniform convergence of $D_N f$ to D_f on $[0,4]$. Thus such diophantine issues are covered up by taking concave hulls.

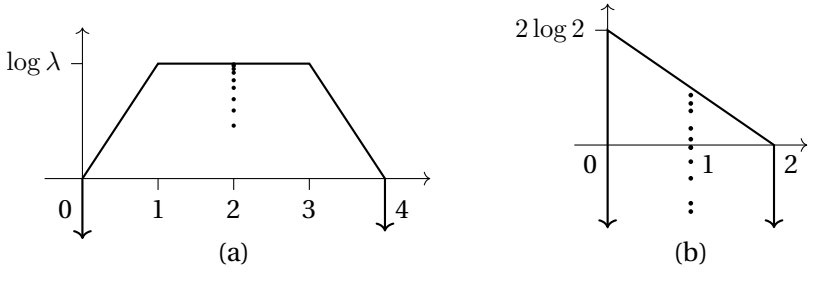

Figure 3.2: (a) Convergence in Example 3.2.2, and (b) lack of convergence in Example 3.2.3

The next example shows that if we allow the coefficients of f to be arbitrary complex numbers instead of integers, then $(L_N f)(k)$ can fail badly to converge at some k.

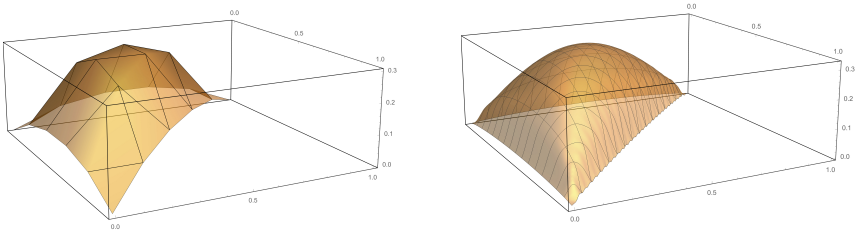

(a) (b)

Figure 3.3: (a) Polyhedral approximation $D_5 f$, and (b) limiting smooth surface D_f for $f(x,y) = 1 + x + y$ in Example 3.2.4

Example 3.2.3. Let $d = 1$ and $f(x) = (x - 2e^{2\pi i\theta})(x - 2e^{-2\pi i\theta})$, where we will determine θ. Then $(L_N f)(0) = 2\log 2$ and $(L_N f)(2) = 0$ for all $N \geqslant 1$, while

$$(L_N f)(1) = \frac{1}{N} \log |2^N \cdot 2\cos(2\pi N\theta)|.$$

It is possible to construct an irrational θ and a sequence $N_j \to \infty$ such that $\frac{1}{N_j} \log |\cos(2\pi N_j\theta)| \to -\infty$ as $j \to \infty$ Hence using this value of θ to define f we see that $(L_N f)(1)$ does not converge, as depicted in Figure 3.2(b), although the concave hulls $D_N f$ do converge uniformly to D_f.

Using arguments similar to those above, it is possible to give an elementary direct proof of Theorem 3.1.1 in the case $d = 1$.

Example 3.2.4. Let $d = 2$ and $f(x,y) = 1 + x + y$. Then f_N is a polynomial in x^N and y^N of degree N^2. For example,

$$\begin{aligned} f_{\langle 5\rangle}(x,y) = {} & x^{25} + 5x^{20}y^5 + 5x^{20} + 10x^{15}y^{10} - 605x^{15}y^5 + 10x^{15} + 10x^{10}y^{15} \\ & + 1905x^{10}y^{10} + 1905x^{10}y^5 + 10x^{10} + 5x^5y^{20} - 605x^5y^{15} + 1905x^5y^{10} \\ & - 605x^5y^5 + 5x^5 + y^{25} + 5y^{20} + 10y^{15} + 10y^{10} + 5y^5 + 1. \end{aligned}$$

The Nth logarithmic rescaling $L_N f$ of f is finite at points in the unit simplex $\Delta = \mathcal{N}_f$ whose coordinates are integer multiples of $1/N$. Thus its concave hull $D_N f$ is a polyhedral surface over Δ, and as $N \to \infty$ these surfaces converge uniformly on Δ to the graph of the concave decimation limit D_f. Figure 3.3(a) shows the polyhedral surface $D_5 f$ corresponding to the calculation of $f_{\langle 5\rangle}$ above, and Figure 3(b) depicts the limiting smooth surface for D_f.

For this example it is possible to derive an explicit formula for D_f. Clearly

$D_f(r,s)$ is symmetric in r and s, so we may assume that $s \leqslant r$. Let

$$\Delta_1 = \{(r,s) \in \Delta : s \leqslant r \text{ and } s \leqslant (1-r)/2\}, \tag{3.6}$$

$$\Delta_2 = \{(r,s) \in \Delta : s \leqslant r \text{ and } s \geqslant (1-r)/2\}. \tag{3.7}$$

For $(r,s) \in \Delta_1 \cup \Delta_2$ with $r+s<1$ define

$$b(r,s) = \csc[\pi(r+s)]\sin(\pi s).$$

Then it turns out that $0 \leqslant b(r,s) \leqslant 1$ for $(r,s) \in \Delta_1$ while $1 \leqslant b(r,s) < \infty$ for $(r,s) \in \Delta_2$.

Using Legendre duality and calculations of R_f by Lundqvist [70], we will show in Appendix A that if $(r,s) \in \Delta_1$ then

$$D_f(r,s) = \sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{\pi n^2} b(r,s)^n \sin[n\pi(1-r)] - s\log b(r,s), \tag{3.8}$$

while if $(r,s) \in \Delta_2$ then

$$D_f(r,s) = \sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{\pi n^2} b(r,s)^{-n} \sin[n\pi(1-r)] + (1-r-s)\log b(r,s). \tag{3.9}$$

We will prove in Corollary 3.1.2 that the maximum value of D_f equals the entropy of α_f, which is the logarithmic Mahler measure $m(f)$ of f defined in (3.2). In this example, the maximum value is attained at $(1/3,1/3)$, which is in both Δ_1 and Δ_2. Either formula therefore applies, and each gives Smyth's calculation [105] that

$$m(1+x+y) = D_f(1/3,1/3) = \frac{3\sqrt{3}}{4\pi}\sum_{n=1}^{\infty}\frac{\chi_3(n)}{n^2} = \frac{3\sqrt{3}}{4\pi}L(2,\chi_3) \approx 0.3230, \tag{3.10}$$

where χ_3 is the nontrivial character of $\mathbb{Z}/3\mathbb{Z}$ and $L(s,\chi_3)$ is the L-function associated with χ_3.

Unlike the previous example, some decimation limits exhibit non-smooth behavior.

Example 3.2.5. Let $d=2$ and $f(x,y) = 5 + x + x^{-1} + y + y^{-1}$. The decimation limit D_f is depicted in Figure 3.4(a). The non-smooth peak at the origin is due to a "hole" in the amoeba of f, as defined in §3.4 and shown in Figure 3.4(b).

As in the previous example, the decimation limit describes the surface tension for a physical model, in this case dimer tilings of the square-octagon graph (see [55, Fig. 3].

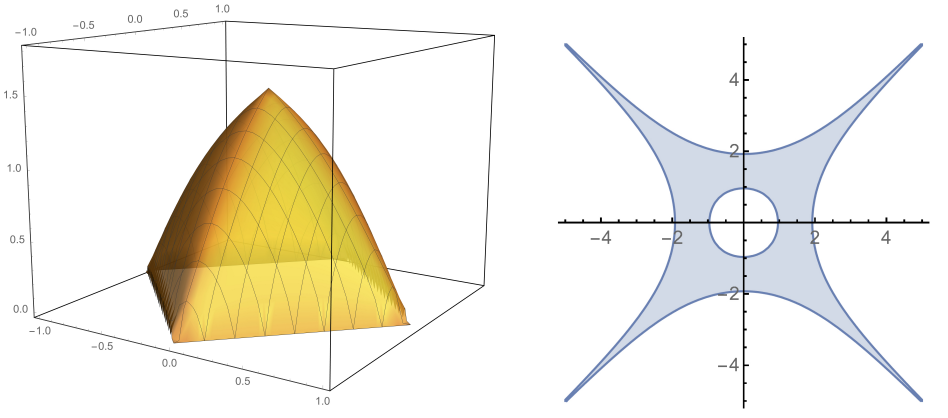

(a) (b)

Figure 3.4: (a) The decimation limit for $f(x,y) = 5 + x + x^{-1} + y + y^{-1}$ from Example 3.2.5 , and (b) the "hole" in its amoeba causing the peak.

Remark 3.2.6. Dimer models have a long history in statistical physics. A particularly important instance involves $f(x,y) = 1 + x + y$ from Example 3.2.4, and has been studied in enormous detail by many authors, including Kenyon, Okounkov, and Sheffield [55].

To describe this model, let $\mathscr{H}$ denote the regular hexagonal lattice in $\mathbb{R}^2$. We can assign the vertices of $\mathscr{H}$ alternating colors red and black, much like a checkerboard. A *perfect matching* on $\mathscr{H}$ is an assignment of each red vertex to a unique adjacent black vertex, these forming an edge or *dimer*. A perfect matching is equivalent to a tiling of $\mathbb{R}^2$ by three types of lozenges, one type for each of the three edges incident to each vertex. Using a natural height function, such a lozenge tiling gives a surface, and the study of the statistical properties of such random surfaces has resulted in many remarkable discoveries (see Kenyon's survey [56], Okounkov's survey [86], or Gorin's detailed account of lozenge tilings [39]).

Kasteleyn discovered that by cleverly assigning signs to the edges of $\mathscr{H}$, he could compute the number of perfect matchings on a finite approximation using periodic boundary conditions by a determinant formula. Furthermore, this determinant can be explicitly evaluated to have the form of a decimation of $f(x,y) = 1 + x + y$. Each of the three terms of f correspond to one of the three types of lozenges in the random tiling. It then turns out that in the logarithmic scaling limit $D_f(r,s)$ counts the growth rate of perfect matchings for which the frequencies of the three lozenge types are r, s, and $1 - r - s$. As such, it is called the *surface tension* for this model.

The two-variable polynomials with integer coefficients arising from such dimer models, such as the preceding two examples, define curves of a very special type called Harnack curves. For these there are probabilistic interpretations of the coefficients of decimations. The additional structure enables one to show that the individual nonzero coefficients of f_N grow at a rate predicted by D_f. Example 3.2.3 shows this can fail if complex coefficients are allowed. But whether or not this is true for every polynomial in $\mathbb{Z}[\mathbb{Z}^d]$ for all $d \geqslant 1$ appears to be quite an interesting problem (see Question 3.9.3 for a precise formulation).

3.3 Convex functions and Legendre duals

We briefly review some basic facts about convex functions and their Legendre duals. Rockafellar's classic book [94] contains a comprehensive account of this theory.

Let $\overline{\mathbb{R}}$ denote $\mathbb{R} \cup \{\infty\}$, with the standard conventions about arithmetic operations and inequalities involving ∞. Let $\varphi\colon \mathbb{R}^d \to \overline{\mathbb{R}}$ be a function, and define its *epigraph* by

$$\operatorname{epi}\varphi := \{(\mathbf{u}, t) : \mathbf{u} \in \mathbb{R}^d, t \in \mathbb{R}, \text{ and } t \geqslant \varphi(\mathbf{u})\} \subset \mathbb{R}^d \times \mathbb{R}.$$

Then φ is *convex* provided that $\operatorname{epi}\varphi$ is a convex subset of $\mathbb{R}^d \times \mathbb{R}$. Similarly, a function $\psi\colon \mathbb{R}^d \to \underline{\mathbb{R}}$ is *concave* if $-\psi\colon \mathbb{R}^d \to \overline{\mathbb{R}}$ is convex.

The *effective domain* of a convex function φ is defined by

$$\operatorname{dom}\varphi := \{\mathbf{u} \in \mathbb{R}^d : \varphi(\mathbf{u}) < \infty\}.$$

By allowing φ to take the value ∞, we may assume that it is defined on all of $\mathbb{R}^d$, enabling us to combine convex functions without needing to take into account their effective domains. A convex function is *closed* if its epigraph is a closed subset of $\mathbb{R}^d \times \mathbb{R}$. This property normalizes the behavior of a convex function at the boundary of its effective domain, and holds for all convex (and concave) functions that arise here.

Suppose that $\varphi\colon \mathbb{R} \to \overline{\mathbb{R}}$ is convex. Its *Legendre dual* (or, more accurately, its *Legendre-Fenchel dual*) φ^* is defined for all $\mathbf{r} \in \mathbb{R}^d$ by

$$\varphi^*(\mathbf{r}) := \sup\{\mathbf{r} \cdot \mathbf{u} - \varphi(\mathbf{u}) : \mathbf{u} \in \mathbb{R}^d\}. \tag{3.11}$$

The Legendre dual φ^* is also a convex function, and provides an alternative description of $\operatorname{epi}\varphi$ in terms of its support hyperplanes. Furthermore, Legendre duality states that $\varphi^{**} = \varphi$ for closed convex functions.

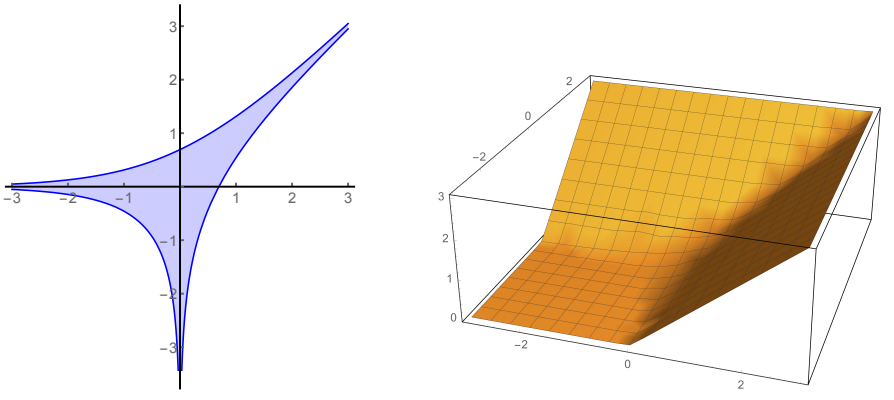

(a) (b)

Figure 3.5: (a) The amoeba of $1 + x + y$, and (b) its Ronkin function

The Legendre dual of a concave function $\psi\colon \mathbb{R}^d \to \underline{\mathbb{R}}$ is similarly defined as

$$\psi^*(\mathbf{r}) = \inf\{\mathbf{r}\cdot\mathbf{u} - \psi(\mathbf{u}) : \mathbf{u} \in \mathbb{R}^d\}. \tag{3.12}$$

Then $\varphi = -\psi$ is convex, and a simple manipulation shows that their Legendre duals are related by $\psi^*(\mathbf{r}) = -\varphi^*(-\mathbf{r})$.

3.4 Amoebas and Ronkin functions

Let $0 \neq f \in \mathbb{C}[\mathbb{Z}^d]$. Put $\mathbb{C}^* = \mathbb{C} \setminus \{0\}$ and define $V(f) := \{\mathbf{z} \in (\mathbb{C}^*)^d : f(\mathbf{z}) = 0\}$. Let $\mathrm{Log}\colon (\mathbb{C}^*)^d \to \mathbb{R}^d$ be the map $\mathrm{Log}(z_1, \ldots, z_d) = (\log|z_1|, \ldots, \log|z_d|)$.

In 1993 Gelfand, Kapranov, and Zelevinsky [36] introduced the notion of the *amoeba* $\mathcal{A}_f$ of f, defined as

$$\mathcal{A}_f := \mathrm{Log}\big(V(f)\big) \subset \mathbb{R}^d.$$

The amoeba of $1 + x + y$ is depicted in Figure 3.5(a). The complement $\mathcal{A}_f^c = \mathbb{R}^d \setminus \mathcal{A}_f$ of $\mathcal{A}_f$ consists of a finite number of connected components, all convex. The unbounded components are created by "tentacles" of $\mathcal{A}_f$. Unfortunately, biological amoebas look nothing like their mathematical namesakes.

Closely related to $\mathcal{A}_f$ is the Ronkin function R_f of f, introduced by Ronkin [95] in 2001, and defined earlier in (3.4). The Ronkin function of $1 + x + y$ is shown in Figure 3.5(b).

The Ronkin function of a polynomial f is known to be a convex function on $\mathbb{R}^d$ and affine on each connected component of $\mathcal{A}_f^c$ (for this and much more see [88]). Moreover, on each connected component of $\mathcal{A}_f^c$ the (constant) gradient of R_f is contained in $\mathcal{N}_f \cap \mathbb{Z}^d$, and the convex hull of these values equals $\mathcal{N}_f$. From this we conclude that the Legendre dual R_f^* of R_f has effective domain $\mathcal{N}_f$.

3.5 Decimation limits of polynomials

In this section we prove Theorem 3.1.1, one of our main results, and Corollaries 3.1.2 and 3.1.3. If $0 \neq f \in \mathbb{C}[\mathbb{Z}^d]$ we will show that the Nth renormalized decimation $D_N f = CH(L_N f)$ converges uniformly on $\mathcal{N}_f$ to a continuous concave limit function D_f, and that $D_f = -R_f^*$.

The first ingredient in our proof is the basic estimate of Mahler relating the largest coefficient of a polynomial to its Mahler measure and its support. Let us begin with some terminology. For $0 \neq g \in \mathbb{C}[\mathbb{Z}^d]$ define its *height* $H(g)$ by $H(g) = \max\{\,|\widehat{g}(\mathbf{k})| : \mathbf{k} \in \mathbb{Z}^d\}$. The *Mahler measure* of g is $M(g) = \exp\big(m(g)\big)$, where $m(g)$ is the logarithmic Mahler measure defined in (3.2).

Proposition 3.5.1 (Mahler [77])**.** Suppose that $0 \neq g \in \mathbb{C}[\mathbb{Z}^d]$ and that $\operatorname{supp} g \subset [0, C-1]^d \cap \mathbb{Z}^d$. Then

$$2^{-dC} H(g) \leqslant M(g) \leqslant C^d H(g). \tag{3.13}$$

Proof. Let $\mathbf{k} = (k_1, \ldots, k_d) \in \operatorname{supp} g$. Then by [77, Eqn. (3)],

$$|\widehat{g}(\mathbf{k})| \leqslant \binom{C-1}{k_1}\binom{C-1}{k_2}\cdots\binom{C-1}{k_d} M(g).$$

Since each binomial coefficient is bounded above by 2^C, the first inequality in (3.13) follows.

The second inequality is a simple consequence of the triangle inequality, since

$$M(g) \leqslant \sum_{\mathbf{k}\in[0,C-1]^d\cap\mathbb{Z}^d} |\widehat{g}(\mathbf{k})| \leqslant |[0,C-1]^d \cap \mathbb{Z}^d| \cdot H(g) = C^d H(g). \qquad \square$$

Consider $(\mathbb{C}^*)^d$ as a group under coordinate-wise multiplication. Define the action of $\mathbf{z} \in (\mathbb{C}^*)^d$ on $f \in \mathbb{C}[\mathbb{Z}^d]$ by $(\mathbf{z}\cdot f)(x_1, \ldots, x_d) = f(z_1 x_1, \ldots, z_d x_d)$. This action is commutative since

$$\mathbf{z}\cdot(\mathbf{z}'\cdot f) = (\mathbf{z}\mathbf{z}')\cdot f = \mathbf{z}'\cdot(\mathbf{z}\cdot f),$$

and also $\mathbf{z}\cdot(fg) = (\mathbf{z}\cdot f)(\mathbf{z}\cdot g)$ for all $f, g \in \mathbb{C}[\mathbb{Z}^d]$. Hence the map $f \mapsto \mathbf{z}\cdot f$ is a ring isomorphism of $\mathbb{C}[\mathbb{Z}^d]$. Furthermore, $(\mathbf{z}\cdot f)\widehat{\ }(\mathbf{k}) = \mathbf{z}^{\mathbf{k}}\widehat{f}(\mathbf{k})$ for all $\mathbf{k} \in \mathbb{Z}^d$, and so $\mathcal{N}_{\mathbf{z}\cdot f} = \mathcal{N}_f$ for all $\mathbf{z} \in (\mathbb{C}^*)^d$.

Recall that Ω_N denotes the group of Nth roots of unity. For $\boldsymbol{\omega} \in \Omega_N^d \subset (\mathbb{C}^*)^d$ we call $\boldsymbol{\omega}\cdot f$ the *rotate of f by $\boldsymbol{\omega}$*. Then $f_N = \prod_{\boldsymbol{\omega}\in\Omega_N^d} \boldsymbol{\omega}\cdot f$ is the product of all rotates of f by elements in Ω_N^d.

If $g, h \in \mathbb{C}[\mathbb{Z}^d]$ then is it well known that $\mathcal{N}_{gh} = \mathcal{N}_f + \mathcal{N}_g$ (the Minkowski sum), and trivially $R_{gh} = R_g + R_h$. By our previous remarks,

$$\mathcal{N}_{f_N} = \sum_{\boldsymbol{\omega}\in\Omega_N^d} \mathcal{N}_{\boldsymbol{\omega}\cdot f} = \sum_{\boldsymbol{\omega}\in\Omega_N^d} \mathcal{N}_f = N^d\mathcal{N}_f.$$

Also, $R_{\boldsymbol{\omega}\cdot f} = R_f$, and hence $R_{f_N} = N^d R_f$.

For $\mathbf{u} \in \mathbb{R}^d$ put $e^{\mathbf{u}} = (e^{u_1}, \dots, e^{u_d})$. Then $(e^{\mathbf{u}}\cdot f)\widehat{\ }(\mathbf{k}) = e^{\mathbf{u}\cdot\mathbf{k}}\widehat{f}(\mathbf{k})$. Commutativity of the action of $(\mathbb{C}^*)^d$ on f then shows that $(e^{\mathbf{u}}\cdot f)_{\langle N\rangle} = e^{\mathbf{u}}\cdot(f_N)$. Also

$$R_f(\mathbf{u}) = \log M(e^{\mathbf{u}}\cdot f) = \frac{1}{N^d}\log M\big((e^{\mathbf{u}}\cdot f)_{\langle N\rangle}\big) = \frac{1}{N^d}\log M(e^{\mathbf{u}}\cdot f_N).$$

Observe that

$$\log H(e^{\mathbf{u}}\cdot f_N) = \max\{\mathbf{u}\cdot\mathbf{k} + \log|\widehat{f_N}(\mathbf{k})| : \mathbf{k} \in \mathbb{Z}^d\},$$

indicating a connection with Legendre duals.

Proof of Theorem 3.1.1. Let $0 \neq f \in \mathbb{C}[\mathbb{Z}^d]$. For $\mathbf{m} \in \mathbb{Z}^d$ let $g(\mathbf{x}) = \mathbf{x}^{\mathbf{m}}f(\mathbf{x})$. It is straightforward to verify that $(D_N g)(\mathbf{r}) = (D_N f)(\mathbf{r} - \mathbf{m})$ for all $\mathbf{r} \in \mathbb{R}^d$. Therefore by adjusting f by suitable monomial, we may assume that $\operatorname{supp} f \subset [0, B-1]^d \cap \mathbb{Z}^d$ for some $B \geqslant 1$. Then $\operatorname{supp}(e^{\mathbf{u}}\cdot f_N) \subset [0, N^d(B-1)]^d \cap \mathbb{Z}^d \subset [0, N^dB - 1]^d \cap \mathbb{Z}^d$ for every $\mathbf{u} \in \mathbb{R}^d$. By Prop.3.5.1,

$$\begin{aligned} R_f(\mathbf{u}) = \frac{1}{N^d}\log M(e^{\mathbf{u}}\cdot f_N) &\leqslant \frac{1}{N^d}\Big\{\log\big[(N^dB)^d\big] + \log H(e^{\mathbf{u}}\cdot f_N)\Big\} \\ &= \frac{\log\big[(N^dB)^d\big]}{N^d} + \frac{1}{N^d}\max_{\mathbf{k}\in\mathbb{Z}^d}\{\mathbf{u}\cdot\mathbf{k} + \log|\widehat{f_N}(\mathbf{k})|\}, \end{aligned}$$

where the error term $b_N := N^{-d}\log\big[(N^dB)^d\big] \to 0$ as $N \to \infty$, uniformly for $\mathbf{u} \in \mathbb{R}^d$.

An opposite inequality is based of the following fundamental observation, used both by Boyd [11] and Purbhoo [92] for different purposes. As we noticed before, f_N is a polynomial in the Nth powers of the variables. Therefore $E_N\widehat{f_N}$ is again a polynomial to which we can apply Prop. 3.5.1, but with

improved constants since the support has now shrunk by a factor of N. This improvement is crucial.

Specifically,

$$\operatorname{supp}(e^{\mathbf{u}} \cdot f_N) \subset [0, N^d(B-1)]^d \cap (N\mathbb{Z}^d),$$

so that

$$\operatorname{supp}(e^{\mathbf{u}} \cdot E_N f_N) \subset [0, N^{d-1}(B-1)] \cap \mathbb{Z}^d.$$

Applying Prop. 3.5.1,

$$H(e^{\mathbf{u}} \cdot f_N) = H(e^{\mathbf{u}} \cdot E_N f_N) \leqslant 2^{dN^{d-1}B} M(e^{\mathbf{u}} \cdot E_N f_N) = 2^{dN^{d-1}B} M(e^{\mathbf{u}} \cdot f)^{N^d}.$$

Hence

$$\frac{1}{N^d} \log H(e^{\mathbf{u}} \cdot f_N) \leqslant \frac{dN^{d-1}B \log 2}{N^d} + \log M(e^{\mathbf{u}} \cdot f) = a_N + R_f(\mathbf{u}),$$

where again the error term $a_N := (dB \log 2)/N \to 0$ uniformly for $\mathbf{u} \in \mathbb{R}^d$. We can summarize these estimates as

$$\Big| R_f(\mathbf{u}) - \frac{1}{N^d} \max_{\mathbf{k} \in \mathbb{Z}^d} \big\{ \mathbf{u} \cdot \mathbf{k} + \log |\widehat{f}_N(\mathbf{k})| \big\} \Big| \leqslant \max\{a_N, b_N\} \to 0 \tag{3.14}$$

as $N \to \infty$ uniformly in $\mathbf{u} \in \mathbb{R}^d$.

Next we relate the first max occurring in (3.14) with the Nth normalized decimation $D_N f$. We have that

$$\begin{aligned}
\frac{1}{N^d} \max_{\mathbf{k} \in \mathbb{Z}^d} \{ \mathbf{u} \cdot \mathbf{k} + \log |\widehat{f}_N(\mathbf{k})| \} &= \max_{\mathbf{k} \in \mathbb{Z}^d} \Big\{ \mathbf{u} \cdot \Big(\frac{\mathbf{k}}{N^d} \Big) + \frac{1}{N^d} \log |\widehat{f}_N(\mathbf{k})| \Big\} \\
&= \max_{\mathbf{k} \in \mathbb{Z}^d} \Big\{ \mathbf{u} \cdot \Big(\frac{\mathbf{k}}{N^d} \Big) + \frac{1}{N^d} E_{N^d} \log \big| \widehat{f}_N \Big(\frac{1}{N^d} \mathbf{k} \Big) \big| \Big\} \\
&= \max_{\mathbf{k} \in \mathbb{Z}^d} \Big\{ \mathbf{u} \cdot \Big(\frac{\mathbf{k}}{N^d} \Big) + (D_N f) \Big(\frac{1}{N^d} \mathbf{k} \Big) \Big\} \\
&= \max_{\mathbf{r} \in \mathbb{R}^d} \{ \mathbf{u} \cdot \mathbf{r} + D_N f(\mathbf{r}) \} = -(D_N f)^*(-\mathbf{u}).
\end{aligned}$$

Hence by (3.14), $-(D_N f)^*(-\mathbf{u})$ converges to $R_f(\mathbf{u})$ uniformly for $\mathbf{u} \in \mathbb{R}^d$, or, equivalently,

$$(D_N f)^*(\mathbf{u}) \to -R_f(-\mathbf{u}) \quad \text{uniformly for } \mathbf{u} \in \mathbb{R}^d. \tag{3.15}$$

If φ and ψ are concave functions on $\mathbb{R}^d$ such that $|\varphi(\mathbf{u}) - \psi(\mathbf{u})| \leqslant \varepsilon$ for all $\mathbf{u} \in \mathbb{R}^d$, it is easy to check from the definitions that φ^* and ψ^* have the same effective domain, and that $|\varphi^*(\mathbf{r}) - \psi^*(\mathbf{r})| \leqslant \varepsilon$ for all $\mathbf{r} \in \operatorname{dom} \varphi^* = \operatorname{dom} \psi^*$. Applying this to (3.15) and using duality we finally obtain that $(D_N f)^{**} = D_N f \to -R_f^*$ uniformly on $\mathcal{N}_f$, completing the proof. □

Proof of Cor. 3.1.2: By Theorem 3.1.1, Legendre duality, and (3.12),

$$-m(f) = -R_f(0,0) = D_f^*(0,0) = \inf_{(r,s)\in\mathcal{N}_f} -D_f(r,s) = -\sup_{(r,s)\in\mathcal{N}_f} D_f(r,s).$$

We remark that differentiability of D_f at the maximum value is not assumed for Legendre duality to apply here, and Example 3.2.5 provides a case when differentiability fails. □

Proof of Cor. 3.1.3: Let f and g be nonzero polynomials in $\mathbb{C}[\mathbb{Z}^d]$. Clearly $R_{fg} = R_f + R_g$. By [94, Thm. 16.4], the Legendre dual of the sum $\varphi + \psi$ of two convex functions is their infimal convolution defined for $\mathbf{r} \in \mathbb{R}^d$ by $\inf\{\varphi(\mathbf{s}) + \psi(\mathbf{r}-\mathbf{s}) : \mathbf{s} \in \mathbb{R}^d\}$. Applying this with $\varphi = -R_f$ and $\psi = -R_g$, using Thm. 3.1.1, and taking negatives we obtain that $D_{fg} = D_f \circledast D_g$. □

Remark 3.5.2. Our estimate (3.14) can be expressed in the language of tropicalization of polynomials (see [74, §3.1] for background and motivation). Let $0 \neq g(\mathbf{x}) = \sum_{\mathbf{k}\in\mathbb{Z}^d} \widehat{g}(\mathbf{k})\mathbf{x}^{\mathbf{k}} \in \mathbb{C}[\mathbb{Z}^d]$. Define the *tropicalization* of g to be the function $\operatorname{trop} g\colon \mathbb{R}^d \to \mathbb{R}$ given by

$$(\operatorname{trop} g)(\mathbf{u}) = \max_{\mathbf{k}\in\mathbb{Z}^d}\big\{\mathbf{u}\cdot\mathbf{k} + \log|\widehat{g}(\mathbf{k})|\big\},$$

which is a polyhedral convex function. Then by (3.14) we see that

$$\frac{1}{N^d}\operatorname{trop} f_N \to R_f \quad \text{uniformly on } \mathbb{R}^d, \tag{3.16}$$

so that the normalized tropicalization of f_N converges uniformly to the Ronkin function of f. Figure 3.6(a) depicts this polyhedral approximation for $f(x,y) = 1+x+y$ and $N = 5$ (compare with Figure 3.5(b)). The tropical variety of this polyhedral approximation is the projection to the plane of the vertices and edges of its graph, and is shown in 3.6(b). These tropical varieties converge in the Hausdorff metric to the amoeba of f as $N \to \infty$ (compare with Figure 3.5(a)).

Remark 3.5.3. In [92] Purbhoo used decimations for a different purpose, namely to find a computational way to detect whether or not a point is in the amoeba of a given polynomial. Call a polynomial *lopsided* if it has one coefficient whose absolute value strictly exceeds the sum of the absolute values of all the other coefficients. Let $f \in \mathbb{Z}[\mathbb{Z}^d]$ and $\mathbf{u} \in \mathbb{R}^d$. Clearly if $e^{\mathbf{u}} \cdot f$ is lopsided then $\mathbf{u} \notin \mathcal{A}_f$. Purbhoo used decimations to amplify size differences among the coefficients. More precisely, he proves that given $\varepsilon > 0$ there is an N_0, depending only on ε and the support of f, such that if the distance from $\mathbf{u}$ to $\mathcal{A}_f$ is greater than ε then $e^{\mathbf{u}}\cdot f_N$ is lopsided. Since f and f_N

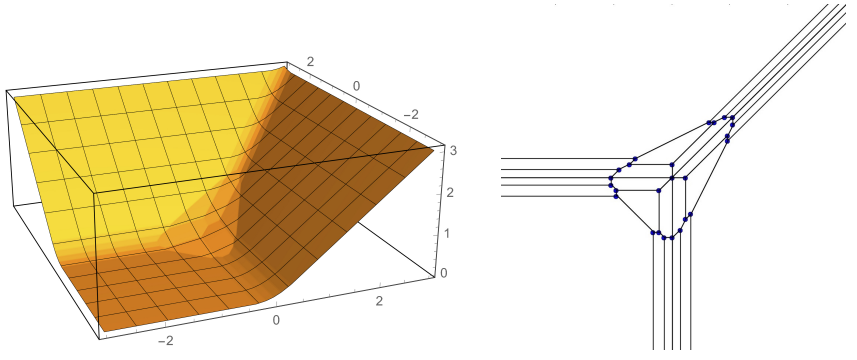

(a) (b)

Figure 3.6: (a) Tropical approximation to the Ronkin function of $1 + x + y$, and (b) its corresponding tropical variety

have the same amoeba, this gives an effective algorithm for approximating the complement of $\mathcal{A}_f$.

One direct consequence of [92] is that the normalized tropicalizations in (3.16) converge to the Ronkin function off the amoeba of f, while our result is that this convergence is uniform on all of $\mathbb{R}^d$. Roughly speaking, Purbhoo is concerned with the coefficients of $e^{\mathbf{u}} \cdot f$ for points $\mathbf{u}$ off the amoeba, while our focus is on $\mathbf{u}$ within the amoeba.

Remark 3.5.4. Let F be a lower-dimensional face of the Newton polytope $\mathcal{N}_f$ of f, and put $f|_F = \sum_{\mathbf{n}\in F} \widehat{f}(\mathbf{n})\mathbf{x}^{\mathbf{n}}$. Clearly the restriction of D_f to F is just the decimation limit of $f|_F$, or in symbols $D_f|_F = D_{f|_F}$. By Corollary 3.1.2, this generalizes [64, Rem. 5.5], which gave a dynamical proof of the inequality due to Smyth [106, Thm. 2] that $m(f) \geqslant m(f_F)$ for every face F of $\mathcal{N}_f$.

3.6 Decimations of principal actions and contracted ideals

We return to decimations of principal algebraic $\mathbb{Z}^d$-actions, and in this section show that they are again principal. The proof uses machinery from commutative algebra, including contractions of ideals.

Suppose that X is a compact, shift-invariant subgroup of $\mathbb{T}^{\mathbb{Z}^d}$. Using Pontryagin duality we can obtain an alternative description of X as follows (for a comprehensive account see [98, Chap. II)].

As a discrete abelian group the Pontryagin dual of $\mathbb{T}^{\mathbb{Z}^d}$ is the direct sum of $\mathbb{Z}^d$ copies of $\mathbb{Z}$, which we suggestively write as $\bigoplus_{\mathbf{k}\in\mathbb{Z}^d}\mathbb{Z}\mathbf{x}^{\mathbf{k}} = \mathbb{Z}[\mathbb{Z}^d]$. The (additive) dual pairing between $\mathbb{T}^{\mathbb{Z}^d}$ and $\mathbb{Z}[\mathbb{Z}^d]$ is given by $\langle t, g\rangle = \sum_{\mathbf{k}\in\mathbb{Z}^d} t_{\mathbf{k}}\,\widehat{g}(\mathbf{k}) \in \mathbb{T}$. Multiplication by the inverses of each of the variables x_j on $\mathbb{Z}[\mathbb{Z}^d]$ gives a $\mathbb{Z}^d$-action that is dual to the natural shift action σ on $\mathbb{T}^{\mathbb{Z}^d}$ defined earlier.

Since X is shift-invariant, $\{g \in \mathbb{Z}[\mathbb{Z}^d] : \langle t, g\rangle = 0 \text{ for all } t \in X\}$ is an ideal $\mathfrak{a}$ in $\mathbb{Z}[\mathbb{Z}^d]$, and the dual group of X equals $\mathbb{Z}[\mathbb{Z}^d]/\mathfrak{a}$. Conversely, if $\mathfrak{a}$ is an arbitrary ideal in $\mathbb{Z}[\mathbb{Z}^d]$, then the compact dual group $X_{\mathfrak{a}}$ of $\mathbb{Z}[\mathbb{Z}^d]/\mathfrak{a}$ is a shift-invariant subgroup of $\mathbb{T}^{\mathbb{Z}^d}$. Thus there is a one-to-one correspondence between shift-invariant compact subgroups of $\mathbb{T}^{\mathbb{Z}^d}$ and ideals in $\mathbb{Z}[\mathbb{Z}^d]$. When $\mathfrak{a}$ is the principal ideal $\langle f\rangle$ generated by f, then $X_{\mathfrak{a}} = X_f$ as defined above, explaining the terminology "principal actions."

Fix $N \geqslant 1$ and recall the restriction map $r_N\colon \mathbb{T}^{\mathbb{Z}^d} \to \mathbb{T}^{N\mathbb{Z}^d}$ from §3.1. Let $f \in \mathbb{Z}[\mathbb{Z}^d]$. Then the Nth decimation $r_N(X_f)$ is a compact subgroup of $\mathbb{T}^{N\mathbb{Z}^d}$ that is invariant under the shift-action of $N\mathbb{Z}^d$. By our previous discussion, the dual group of $r_N(X_f)$ has the form $\mathbb{Z}[N\mathbb{Z}^d]/\mathfrak{a}_N$, where $\mathfrak{a}_N$ is an ideal in $\mathbb{Z}[N\mathbb{Z}^d]$. The following result identifies this ideal.

Lemma 3.6.1. Let $f \in \mathbb{Z}[\mathbb{Z}^d]$ and $N \geqslant 1$. Then the dual group of $r_N(X_f)$ is $\mathbb{Z}[N\mathbb{Z}^d]/\mathfrak{a}_N$, where $\mathfrak{a}_N = \langle f\rangle \cap \mathbb{Z}[N\mathbb{Z}^d]$.

Proof. Let $\mathfrak{b}_N = \{g \in \mathbb{Z}[N\mathbb{Z}^d] : \langle t, g\rangle = 0 \text{ for all } t \in r_N(X_f)\}$. If $g \in \mathfrak{a}_N$, then for every $t \in X_f$ we have that $0 = \langle t, g\rangle = \langle r_N(t), g\rangle$, so that $g \in \mathfrak{b}_N$. Conversely, if $g \in \mathfrak{b}_N$ and $t \in X_f$, then g annihilates the restriction of t to every coset of $N\mathbb{Z}^d$, and hence annihilates t, so that $g \in \mathfrak{a}_N$. □

The ideal $\langle f\rangle \cap \mathbb{Z}[N\mathbb{Z}^d]$ defining $r_N(X_f)$ is called the *contraction* of $\langle f\rangle$ to $\mathbb{Z}[N\mathbb{Z}^d]$. The main result of this section is that this contraction is always principal.

Proposition 3.6.2. Let $f \in \mathbb{Z}[\mathbb{Z}^d]$ and $N \geqslant 1$. Then the contracted ideal $\langle f\rangle \cap \mathbb{Z}[N\mathbb{Z}^d]$ is a principal ideal in $\mathbb{Z}[N\mathbb{Z}^d]$.

We begin by briefly sketching the necessary terminology and machinery from commutative algebra, all of which is contained in [5] or can be easily deduced from material there.

For brevity let $R = \mathbb{Z}[N\mathbb{Z}^d]$ and $S = \mathbb{Z}[\mathbb{Z}^d]$. Both R and S are unique factorization domains, and therefore both are integrally closed [5, Prop. 5.12]. Furthermore, S is integral over R since each variable x_j in S satisfies the monic polynomial $y^N - x_j^N \in R[y]$.

A prime ideal $\mathfrak{p}$ in an integral domain has *height one* if there are no prime ideals strictly between 0 and $\mathfrak{p}$. In a unique factorization domain the prime ideals of height one are exactly the principal ideals generated by irreducible elements. A proper ideal $\mathfrak{q}$ in an integral domain is *primary* if whenever $ab \in \mathfrak{q}$ then either $a \in \mathfrak{q}$ or $b^n \in \mathfrak{q}$ for some $n \geqslant 1$. In this case its radical $\{a : a^n \in \mathfrak{q} \text{ for some } n \geqslant 1\}$ is a prime ideal, say $\mathfrak{p}$, and then $\mathfrak{q}$ is called $\mathfrak{p}$-*primary*. Examples show that in general a power of a prime ideal need not be primary, that a primary ideal need not be the power of a prime ideal, and that even if an ideal has prime radical it need not be primary. The notion of primary ideal, although the correct one for decomposition theory, is quite subtle. However, in our situation things are much simpler.

Lemma 3.6.3. Let P be a unique factorization domain, and let $r \in P$ be irreducible. Then the principal ideal $\mathfrak{p} = \langle r \rangle$ is prime, and the $\mathfrak{p}$-primary ideals are exactly the powers $\mathfrak{p}^n$ of $\mathfrak{p}$ for $n \geqslant 1$.

Proof. It is clear that $\mathfrak{p}$ is prime. To prove that $\mathfrak{p}^n = \langle r^n \rangle$ is $\mathfrak{p}$-primary, suppose that $ab \in \mathfrak{p}^n$, but $a \notin \mathfrak{p}^n$. Then $r \mid b$, so $b^n \in \mathfrak{p}^n$, showing that $\mathfrak{p}^n$ is primary. Clearly the radical of $\mathfrak{p}^n$ is $\mathfrak{p}$, and so $\mathfrak{p}^n$ is $\mathfrak{p}$-primary.

Conversely, suppose that $\mathfrak{q}$ is a $\mathfrak{p}$-primary ideal. Since the radical of $\mathfrak{q}$ is $\mathfrak{p}$, it follows that $r^n \in \mathfrak{q}$ for some $n \geqslant 1$. Choose n to be the minimal such power. Then $\mathfrak{p}^n \subset \mathfrak{q}$, but there is an $a \in \mathfrak{q} \smallsetminus \mathfrak{p}^{n-1}$. Write $a = cr^m$, where $r \nmid c$. Choose a so that m is the maximal such power, where obviously $m \leqslant n-1$. Now $r^n \notin \mathfrak{q}$ by minimality of n, hence some power $c^k \in \mathfrak{q} \subset \mathfrak{p}$ since $\mathfrak{q}$ is primary. But this is absurd since $r \nmid c$ unless c is a unit. Thus $\mathfrak{q} = \mathfrak{p}^{n-1}$. □

If $\mathfrak{a}$ is an ideal in S, we denote its *contraction* $\mathfrak{a} \cap R$ to R by $\mathfrak{a}^c$. If $\mathfrak{q}$ is a $\mathfrak{p}$-primary ideal in S, then $\mathfrak{p}^c$ is prime and $\mathfrak{q}^c$ is $\mathfrak{p}^c$-primary in R.

One of the important results in commutative algebra, essential to developing a dimension theory using chains of prime ideals, is the so-called "Going Down" theorem [5, Thm. 5.16]. Its hypotheses are satisfied in our situation, and it says the following. Suppose that $\mathfrak{p}_0 \subsetneq \mathfrak{p}_1 \subsetneq \mathfrak{p}_2$ is a chain of prime ideals in R, and that there is a prime ideal $\mathfrak{q}_2$ in S with $\mathfrak{q}_2^c = \mathfrak{p}_2$. Then there is a chain $\mathfrak{q}_0 \subsetneq \mathfrak{q}_1 \subsetneq \mathfrak{q}_2$ of prime ideals in S such that $\mathfrak{q}_j^c = \mathfrak{p}_j$ for $j = 0, 1, 2$. From this it follows that prime ideals in S of height one contract to prime ideals in R of height one. In other words, if $h \in S$ is irreducible, then $\langle h \rangle_S \cap R$ is a principal ideal $\langle g \rangle_R$ in R generated by an irreducible polynomial g in R.

Proof of Prop. 3.6.2. First suppose that $f \in S$ is irreducible. As we just showed, there is an irreducible $g \in R$ such that $\langle f \rangle_S \cap R = \langle g \rangle_R$. Furthermore, if $n \geqslant 1$

then $\langle f^n\rangle_S$ is $\langle f\rangle_S$-primary, and so $\langle f^n\rangle_S \cap R$ is $\langle g\rangle_R$-primary, hence equals $\langle g^k\rangle_R$ for some $k \geqslant 1$.

The result is obvious if $f = 0$, so suppose that $0 \neq f \in S$, and let $f = f_1^{n_1} \cdots f_r^{n_r}$ be its factorization in S into powers of distinct irreducibles f_j. Then there are irreducible polynomials $g_j \in R$ and $k_j \geqslant 1$ such that $\langle f_j^{n_j}\rangle_S \cap R = \langle g_j^{k_j}\rangle_R$. Hence

$$\begin{aligned}\langle f\rangle_S \cap R &= \langle f_1^{n_1} \cdots f_r^{n_r}\rangle_S \cap R = \big(\langle f_1^{n_1}\rangle_S \cap \cdots \langle f_r^{n_r}\rangle_S\big) \cap R \\ &= \big(\langle f_1^{n_1}\rangle_S \cap R\big) \cap \cdots \cap \big(\langle f_r^{n_r}\rangle_S \cap R\big) \\ &= \langle g_1^{k_1}\rangle_R \cap \cdots \cap \langle g_r^{k_r}\rangle_R = \langle \mathrm{LCM}(g_1^{k_1}, \ldots, g_r^{k_r})\rangle_R,\end{aligned}$$

proving that $\langle f\rangle_S \cap R$ is principal. □

Remarks 3.6.4. (1) It is possible for distinct principal prime ideals in S to contract to the same prime ideal in R. As a simple example, let $d = 1$, $N = 2$, $f_1(x) = x^2 - x - 1$, and $f_2(x) = x^2 + x - 1$. Then each is irreducible in S, but both $\langle f_1\rangle_S$ and $\langle f_2\rangle_S$ contract in $R = \mathbb{Z}[2\mathbb{Z}]$ to $\langle x^4 - 3x^2 + 1\rangle_R$, where $x^4 - 3x^2 + 1$ is irreducible in $\mathbb{Z}[2\mathbb{Z}]$ (but of course not in $\mathbb{Z}[\mathbb{Z}]$). In the proof this is accounted for by using the least common multiple LCM in the last line of the displayed equation above.

(2) A polynomial is *primitive* if the greatest common divisor of its coefficients is 1. If $0 \neq f \in S$ is a nonconstant primitive polynomial with factorization $f = f_1^{n_1} \cdots f_r^{n_r}$ into powers of distinct irreducible polynomials, then by Gauss's Lemma each f_j is primitive as well. Furthermore, $\langle f_j\rangle_S \cap R = \langle g_j\rangle_R$, where each g_j is nonconstant and primitive. It then follows from the proof that $\langle f\rangle_S \cap R$ is generated by a primitive element of R.

(3) There is a completely different proof of Prop. 3.6.2 using entropy that is valid for all polynomials in S except for those of a very special and easily determined form. Recall that the entropy of α_f is the logarithmic Mahler measure $m(f)$ defined in (3.2). A *generalized cyclotomic polynomial* in S is one of the form $\mathbf{x}^{\mathbf{n}}c(\mathbf{x}^{\mathbf{k}})$, where c is a cyclotomic polynomial in one variable and $\mathbf{k} \neq \mathbf{0}$. Smyth [106] proved that $m(f) = 0$ if and only if f is, up to sign, a product generalized cyclotomic polynomials. Assume that $f \in S$ is not such a polynomial, so that the entropy of α_f is strictly positive. A simple argument using cosets of $N\mathbb{Z}^d$ shows that $r_N(X_f)$ also has positive entropy. Now $r_N(X_f) = X_{\mathfrak{a}_N}$ by Lemma 3.6.1, where $\mathfrak{a}_N = \langle f\rangle_S \cap R$. But an ideal $\mathfrak{a}$ in R for which the shift action of $N\mathbb{Z}^d$ on $X_{\mathfrak{a}}$ has positive entropy must be principal [64, Thm. 4.2].

3.7 Absolutely irreducible factorizations and Gauss's Lemma

Suppose that $f \in \mathbb{Z}[\mathbb{Z}^d]$ is nonconstant and irreducible. Its factorization into absolutely irreducible polynomials in an extension field of $\mathbb{Q}$ will play a decisive role. A generalization of Gauss's Lemma to number fields enables us to deal with the algebraic properties of the coefficients of the factors.

Two polynomials in $\mathbb{C}[\mathbb{Z}^d]$ are *distinct* if one is not a nonzero scalar multiple of the other. An element $\varphi \in \mathbb{C}[\mathbb{Z}^d]$ is *adjusted* if $\mathbf{0}$ is an extreme point of its Newton polytope $\mathcal{N}_\varphi$, and is *monic* if it is both adjusted and $\widehat{\varphi}(\mathbf{0}) = 1$.

A polynomial in $\mathbb{C}[\mathbb{Z}^d]$ is *absolutely irreducible* if it is irreducible in the unique factorization domain $\mathbb{C}[\mathbb{Z}^d]$. Hence every non-unit $f \in \mathbb{C}[\mathbb{Z}^d]$ has some factorization $f = \varphi_1 \cdots \varphi_r$ into absolutely irreducible factors φ_j. The method of Galois descent [20] shows that, after multiplying the factors by suitable constants, there is a finite normal extension $\mathbb{K}$ of $\mathbb{Q}$ such that each $\varphi_j \in \mathbb{K}[\mathbb{Z}^d]$, and also that the coefficients of the φ_j generate $\mathbb{K}$, so that $\mathbb{K}$ is the splitting field of f. Furthermore an elementary argument shows that if f is adjusted, then we can multiply the φ_j by units in $\mathbb{K}[\mathbb{Z}^d]$ so that each φ_j is monic, $\mathcal{N}_{\varphi_j} \subset \mathcal{N}_f$, and $f = \widehat{f}(\mathbf{0})\varphi_1 \cdots \varphi_r$.

Remarks 3.7.1. (1) When $d = 1$ this factorization is into the linear factors guaranteed by the fundamental theorem of algebra.

(2) A simple sufficient condition for φ to be absolutely irreducible is that $\mathcal{N}_\varphi$ is not the nontrivial Minkowski sum of two integer polytopes (see [35] for applications of this idea).

(3) There are reasonably good factoring algorithms which, on input f, produce a monic irreducible polynomial in $\mathbb{Z}[x]$ with root θ and an absolutely irreducible $\varphi \in \mathbb{Q}(\theta)[\mathbb{Z}^d]$ such that $f = \sigma_1(\varphi)\sigma_2(\varphi)\cdots\sigma_r(\varphi)$, where the σ_j are all the distinct field embeddings of $\mathbb{Q}(\theta)$ into $\mathbb{C}$ (see [29] for an overview of these methods).

The following shows that, unlike factoring, divisibility is not affected when passing to an extension field.

Lemma 3.7.2. Suppose that $\mathbb{L}$ is an extension of the field $\mathbb{K}$ and that $f, g \in \mathbb{K}[\mathbb{Z}^d]$. Then f divides g in $\mathbb{K}[\mathbb{Z}^d]$ if and only if f divides g in $\mathbb{L}[\mathbb{Z}^d]$.

Proof. For the nontrivial direction, suppose there is an $h \in \mathbb{L}[\mathbb{Z}^d]$ such that $fh = g$. Equating coefficients of like monomials gives a system of $\mathbb{K}$-linear equations in the coefficients of h. Since this system has a solution over $\mathbb{L}$,

Gaussian elimination shows that that this (unique) solution is actually over $\mathbb{K}$, and so $h \in \mathbb{K}[\mathbb{Z}^d]$. □

Proposition 3.7.3. Let $f \in \mathbb{Z}[\mathbb{Z}^d]$ be nonconstant, adjusted, and irreducible in $\mathbb{Z}[\mathbb{Z}^d]$. Then there is a finite normal extension field $\mathbb{K}$ of $\mathbb{Q}$ and monic absolutely irreducible polynomials $\varphi_1, \ldots, \varphi_r \in \mathbb{K}[\mathbb{Z}^d]$ such that $f = \widehat{f}(\mathbf{0})\varphi_1 \cdots \varphi_r$ and $\mathcal{N}_{\varphi_f} \subset \mathcal{N}_f$ for $1 \leqslant j \leqslant r$. Furthermore, the Galois group $\mathrm{Gal}(\mathbb{K} : \mathbb{Q})$ acts transitively on the set of factors φ_j, and these factors are pairwise distinct.

Proof. Our earlier discussion shows there is a factorization $f = \widehat{f}(\mathbf{0})\varphi_1 \cdots \varphi_r$ over the splitting field $\mathbb{K}$ of f, where each φ_j is monic and $\mathcal{N}_{\varphi_j} \subset \mathcal{N}_f$ for $1 \leqslant j \leqslant r$. Suppose that $\sigma \in \mathrm{Gal}(\mathbb{K} : \mathbb{Q})$. Since $\sigma(f) = f$, it follows that σ must permute the absolutely irreducible factors up to multiplication by units. But if $\sigma(\varphi_j) = c\mathbf{x}^{\mathbf{n}}\varphi_k$, then $\mathbf{n} = \mathbf{0}$ since the factors are adjusted and $c = 1$ since they are monic. Hence σ permutes the factors themselves. If there were a proper subset of factors that is invariant under $\mathrm{Gal}(\mathbb{K} : \mathbb{Q})$, then their product ψ would be in $\mathbb{Q}[\mathbb{Z}^d]$ since its coefficients are invariant under $\mathrm{Gal}(\mathbb{K} : \mathbb{Q})$. But then ψ would be a proper divisor of f in $\mathbb{Q}[\mathbb{Z}^d]$ by Lemma 3.7.2, contradicting irreducibility of f by Gauss's Lemma. A similar argument shows that each factor appears with multiplicity one. □

We now give a brief sketch of the extension of Gauss's Lemma to number fields and the consequences we use. Let $\mathbb{K}$ be a finite extension of $\mathbb{Q}$, and $\mathcal{O}_{\mathbb{K}}$ be the ring of algebraic integers in $\mathbb{K}$. A *fractional ideal* $\mathfrak{a}$ in $\mathbb{K}$ is a nonzero $\mathcal{O}_{\mathbb{K}}$-sumbodule such that there is an integer b for which $b\mathfrak{a} \subset \mathcal{O}_{\mathbb{K}}$. Fractional ideals can be added and multiplied, with $\mathcal{O}_{\mathbb{K}}$ being the multiplicative identity. A fractional ideal contained in $\mathcal{O}_{\mathbb{K}}$ is an ideal in the usual ring-theoretic sense. The pivotal result is that the set of fractional ideals form a group, the set of principal fractional ideals (those of the form $\mathcal{O}_{\mathbb{K}}\beta$ for some $\beta \in \mathbb{K}$) form a subgroup, and the quotient of these groups is a finite abelian group called the class group which measures how far $\mathcal{O}_{\mathbb{K}}$ is from being a principal ideal domain.

Let $\varphi \in \mathbb{K}[\mathbb{Z}^d]$. Define the *content* $\mathfrak{c}_{\mathbb{K}}(\varphi)$ to be the fractional ideal in $\mathbb{K}$ generated by the coefficients of φ. Say that φ is *primitive* if $\mathfrak{c}_{\mathbb{K}}(\varphi) = \mathcal{O}_{\mathbb{K}}$. It is easy to check that although content depends on the ambient field $\mathbb{K}$, primitivity does not: if $\varphi \in \mathbb{K}[\mathbb{Z}^d]$ and $\varphi \in \mathbb{L}[\mathbb{Z}^d]$, then $\mathfrak{c}_{\mathbb{K}}(\varphi) = \mathcal{O}_{\mathbb{K}}$ if and only if $\mathfrak{c}_{\mathbb{L}}(\varphi) = \mathcal{O}_{\mathbb{L}}$ (see [75, Thm. 8.2]).

Theorem 3.7.4 (Gauss's Lemma for number fields)**.** Let $\mathbb{K}$ be a number field and $\varphi, \psi \in \mathbb{K}[\mathbb{Z}^d]$. The $\mathfrak{c}_{\mathbb{K}}(\varphi\psi) = \mathfrak{c}_{\mathbb{K}}(\varphi)\mathfrak{c}_{\mathbb{K}}(\psi)$. In particular, if $\varphi, \psi \in \mathcal{O}_{\mathbb{K}}[\mathbb{Z}^d]$ then $\varphi\psi$ is primitive if and only if both φ and ψ are primitive. If $\varphi, \psi \in \mathcal{O}_{\mathbb{K}}[\mathbb{Z}^d]$ are primitive, and if $\varphi = \beta\psi$ for some $\beta \in \mathbb{K}$, then β is a unit in $\mathcal{O}_{\mathbb{K}}$.

Remark 3.7.5. Suppose that $f \in \mathbb{Z}[\mathbb{Z}^d]$ is primitive and that $N \geqslant 1$. Let $\zeta_N = e^{2\pi i/N}$, which is a unit in $\mathbb{Q}(\zeta_N)$. Hence each rotate $\boldsymbol{\omega} \cdot f$, where $\boldsymbol{\omega} \in \Omega_N^d$, is primitive in $\mathbb{Q}(\zeta_N)[\mathbb{Z}^d]$. The preceding theorem then shows that the product f_N of these rotates is also primitive in $\mathbb{Q}(\zeta_N)[\mathbb{Z}^d]$, and hence in $\mathbb{Z}[\mathbb{Z}^d]$ (since primitivity is independent of ambient field), a fact we already observed in Remark 3.6.4(2).

3.8 Decimated polynomials and decimated actions

Let $f \in \mathbb{Z}[\mathbb{Z}^d]$ be irreducible. Here we explain the relationship between the Nth decimation f_N of f and the generator g_N of the contracted ideal $\langle f \rangle \cap \mathbb{Z}[N\mathbb{Z}^d]$ that defines the Nth decimation $r_N(X_f)$ of (X_f, α_f). Roughly speaking, g_N is a constant times the product of all distinct rotates by elements of Ω_N^d of the absolutely irreducible factors φ_j of f as described in Proposition 3.7.3. Each rotate appears with the same multiplicity e_N that can be computed from the φ_j. Thus $f_N = c\, g_N^{e_N}$, and an application of Gauss's Lemma shows that we may take $c = 1$. Furthermore, there is an integer $Q(f)$, that can also be computed from the φ_j, such that $f_N = g_N$ for all N relatively prime to $Q(f)$. Examples will illustrate the two sources of the multiplicity e_N.

In what follows we let $\zeta_N = e^{2\pi i/N}$, which is a generator of Ω_N.

Lemma 3.8.1. If $f \in \mathbb{Z}[\mathbb{Z}^d]$ then $f_N \in \mathbb{Z}[N\mathbb{Z}^d]$.

Proof. Since $f_N = \prod_{\boldsymbol{\omega} \in \Omega_N^d} \boldsymbol{\omega} \cdot f$, it follows that $f_N = \boldsymbol{\omega} \cdot f_N$ for every $\boldsymbol{\omega} \in \Omega_N^d$. Suppose that $\widehat{f_N}(\mathbf{k}) \neq 0$. Then since

$$\widehat{f_N}(\mathbf{k}) = (\boldsymbol{\omega} \cdot f_N)\widehat{\ }(\mathbf{k}) = \boldsymbol{\omega}^{\mathbf{k}}\, \widehat{f_N}(\mathbf{k}),$$

we see that $\boldsymbol{\omega}^{\mathbf{k}} = 1$ for every $\boldsymbol{\omega} \in \Omega_N^d$, and hence $\mathbf{k} \in N\mathbb{Z}^d$. Thus $f_N \in \mathbb{Q}(\zeta_N)[N\mathbb{Z}^d]$.

The Galois group $G := \mathrm{Gal}\big(\mathbb{Q}(\zeta_N) : \mathbb{Q}\big)$ acts on Ω_N^d coordinate-wise. If $\sigma \in G$, then $\sigma(\boldsymbol{\omega} \cdot f) = \sigma(\boldsymbol{\omega}) \cdot f$ since f has integer coefficients. Thus σ permutes the rotates of f, and so $\sigma(f_N) = f_N$ for every $\sigma \in G$. It follows that the coefficients of f_N are both rational and algebraic integers, and so $f_N \in \mathbb{Z}[N\mathbb{Z}^d]$. □

Lemma 3.8.2. Let $f \in \mathbb{Z}[\mathbb{Z}^d]$ and g_N be a generator of the contracted ideal $\langle f \rangle \cap \mathbb{Z}[N\mathbb{Z}^d]$. Then g_N divides f_N in $\mathbb{Z}[N\mathbb{Z}^d]$.

Proof. Since f is one of the factors in forming f_N, it follows that f divides f_N in $\mathbb{Q}(\zeta_N)[\mathbb{Z}^d]$. Hence f divides f_N in $\mathbb{Q}[\mathbb{Z}^d]$ by Lemma 3.7.2. The coefficients of f_N/f are both rational and algebraic integers, and so $f_N/f \in \mathbb{Z}[\mathbb{Z}^d]$. Hence $f_N \in \langle f \rangle \cap \mathbb{Z}[N\mathbb{Z}^d]$, and it is thus divisible by the generator g_N. □

Remark 3.8.3. Since the generator of a principal ideal is unique only up to units, it will be convenient to have a convention to pick a generator. In what follows we will assume that f is adjusted and that $\widehat{f}(\mathbf{0}) > 0$. Then clearly f_N has the same properties. By the previous lemma, we can also assume that g_N is adjusted, that $\mathcal{N}_{g_N} \subset \mathcal{N}_{f_N}$, and that $\widehat{g}_N(\mathbf{0}) > 0$.

Before continuing, we remark that if f is a constant integer n, then $f_N = n^{N^d}$ while $g_N = n$. Let us call a polynomial $f \in \mathbb{Z}[\mathbb{Z}^d]$ *nonconstant* if $|\operatorname{supp} f| > 1$, and it is these we now turn to.

Let $f \in \mathbb{Z}[\mathbb{Z}^d]$ be adjusted. Define its *support group* Γ_f to be the subgroup of $\mathbb{Z}^d$ generated by $\operatorname{supp} f$. It is easy to check that the support group is independent of which extreme point of $\mathcal{N}_f$ is used to adjust f. We say that f is *full* if $\Gamma_f = \mathbb{Z}^d$.

The following shows that in some cases, including $f(x,y) = 1 + x + y$ from Example 3.2.4, $f_N = g_N$ for all $N \geqslant 1$.

Proposition 3.8.4. Let $f \in \mathbb{Z}[\mathbb{Z}^d]$ be adjusted, irreducible, and full. Further assume that f is absolutely irreducible in $\mathbb{C}[\mathbb{Z}^d]$. Then $f_N = g_N$ for every $N \geqslant 1$.

Proof. Since the map $f \mapsto \boldsymbol{\omega} \cdot f$ is a ring isomorphism of $\mathbb{C}[\mathbb{Z}^d]$, each $\boldsymbol{\omega} \cdot f$ is absolutely irreducible. Suppose that $\boldsymbol{\omega} \cdot f = \boldsymbol{\omega}' \cdot f$. Since $\boldsymbol{\omega}^{\mathbf{k}} \widehat{f}(\mathbf{k}) = (\boldsymbol{\omega}')^{\mathbf{k}} \widehat{f}(\mathbf{k})$, it follow that $\boldsymbol{\omega}^{\mathbf{k}} = (\boldsymbol{\omega}')^{\mathbf{k}}$ for all $\mathbf{k} \in \operatorname{supp} f$, hence for all $\mathbf{k} \in \Gamma_f = \mathbb{Z}^d$, and so $\boldsymbol{\omega} = \boldsymbol{\omega}'$. Thus the rotates of $\boldsymbol{\omega} \cdot f$ for $\boldsymbol{\omega} \in \Omega_N^d$ are pairwise distinct absolutely irreducible polynomials in $\mathbb{C}[\mathbb{Z}^d]$ whose product is f_N.

By Lemma 3.8.2, g_N divides f_N in $\mathbb{C}[\mathbb{Z}^d]$. Hence some rotate $\boldsymbol{\omega} \cdot f$ divides g_N. Since $g_N \in \mathbb{Z}[N\mathbb{Z}^d]$, it is invariant under all rotations in Ω_N^d. Hence g_N is divisible by all rotates $\boldsymbol{\omega} \cdot f$, and so g_N and f_N have the same absolute factorizations in $\mathbb{C}[\mathbb{Z}^d]$, and hence $f_N = c\, g_N$ for some constant $c \in \mathbb{C}$. Recalling our conventions in Remark 3.8.3, comparing constant terms shows that $c = \widehat{f}(\mathbf{0})^{N^d} / \widehat{g}_N(\mathbf{0}) \in \mathbb{Q}$. But f_N and g_N are both primitive in $\mathbb{Z}[N\mathbb{Z}^d]$, and so $c = \pm 1$, and our convention on positivity of constant terms then gives $c = 1$. □

The following example shows that when the polynomial is not full there can be multiplicity $e_N > 1$.

Example 3.8.5. Let $d = 2$ and $f(x, y) = 1+x+y^2$. Since $\mathcal{N}_f$ is not a nontrivial Minkowski sum of integer polytopes, we see that f is absolutely irreducible. Suppose that N is odd. Since $-1 \notin \Omega_N$, all rotates $\boldsymbol{\omega} \cdot f$ for $\boldsymbol{\omega} \in \Omega_N^2$ are distinct, and the same arguments as in the previous proposition show that $f_N = g_N$.

However, if N is even, then $-1 \in \Omega_N$ and the rotate of f by (ω_1, ω_2) equals that by $(\omega_1, -\omega_2)$. As we will see in Proposition 3.8.7, the product of the distinct rotates of f equals g_N, and so $f_N = g_N^2$ when N is even.

Next we characterize when rotates can coincide.

Lemma 3.8.6. Let $\varphi \in \mathbb{C}[\mathbb{Z}^d]$ be adjusted, and Γ_φ be its support group. Then the dual of the stabilizer group $S_N(\varphi) := \{\boldsymbol{\omega} \in \Omega_N^d : \boldsymbol{\omega} \cdot \varphi = \varphi\}$ is $\mathbb{Z}^d/(\Gamma_\varphi + N\mathbb{Z}^d)$. Two rotates of φ differ by a multiplicative unit in $\mathbb{C}[\mathbb{Z}^d]$ if and only if they are equal. If Γ_φ has finite index K in $\mathbb{Z}^d$, then $S_N(\varphi)$ is trivial for every N relatively prime to K.

Proof. Suppose that $\boldsymbol{\omega} \in S_N(\varphi)$. Since $\widehat{\varphi}(\mathbf{k}) = (\boldsymbol{\omega} \cdot \varphi)\widehat{\ }(\mathbf{k}) = \boldsymbol{\omega}^{\mathbf{k}}\widehat{\varphi}(\mathbf{k})$, it follows that $\boldsymbol{\omega}^{\mathbf{k}} = 1$ for every $\mathbf{k} \in \operatorname{supp}\varphi$. Hence $\boldsymbol{\omega}$ annihilates Γ_φ as well as $N\mathbb{Z}^d$, thus their sum. Conversely, every $\boldsymbol{\omega}$ annihilating $\Gamma_\varphi + N\mathbb{Z}^d$ must be in $S_N(\varphi)$. Hence the annihilator of $S_N(\varphi)$ equals $\Gamma_\varphi + N\mathbb{Z}^d$, and so its dual group is $\mathbb{Z}^d/(\Gamma_\varphi + N\mathbb{Z}^d)$.

The multiplicative units in $\mathbb{C}[\mathbb{Z}^d]$ have the form $c\,\mathbf{x}^{\mathbf{n}}$ for some $c \in \mathbb{C}$, so the second statement is obvious since φ is adjusted.

Suppose that Γ_φ has finite index K in $\mathbb{Z}^d$. If N is relatively prime to K, then multiplication by N on $\mathbb{Z}^d/\Gamma_\varphi$ is injective, hence surjective. Thus modulo Γ_φ every element in $\mathbb{Z}^d$ is a multiple of N, and hence $\Gamma_\varphi + N\mathbb{Z}^d = \mathbb{Z}^d$. □

Proposition 3.8.7. Let $f \in \mathbb{Z}[\mathbb{Z}^d]$ be adjusted and irreducible, and further assume that f is absolutely irreducible in $\mathbb{C}[\mathbb{Z}^d]$. Then $f_N = g_N^{e_N}$, where $e_N = |S_N(f)| = |\mathbb{Z}^d/(\Gamma_f + N\mathbb{Z}^d)|$.

Proof. Recall our conventions in Remark 3.8.3. Since g_N divides f_N, it must be divisible by at least one (absolutely irreducible) rotate of f. Invariance of g_N by every rotate in Ω_N^d shows that g_N is therefore divisible by the product h of all the distinct rotates of f. The arguments in Lemmas 3.8.1 and 3.8.2 apply to show that $h \in \langle f \rangle \cap \mathbb{Z}[N\mathbb{Z}^d]$. Thus g_N divides h in $\mathbb{C}[\mathbb{Z}^d]$ as well, and so $g_N = c\,h$ for some $c \in \mathbb{C}$. Evaluating constant terms shows that $c \in \mathbb{Q}$. Since g_N is irreducible in $\mathbb{Z}[N\mathbb{Z}^d]$, it is primitive. Each rotate of f is primitive in $\mathbb{Q}(\zeta_N)[\mathbb{Z}^d]$, and so h is primitive by Theorem 3.7.4. Hence $c = \pm 1$, and

then $c = 1$ follows from our sign conventions. By Lemma 3.8.6, each rotate of f is repeated exactly e_N times, and so $f_N = g_N^{e_N}$. □

When f is absolutely irreducible, the only source of multiplicity $e_N > 1$ is its support group. However, if f has several absolutely irreducible factors, a new source of multiplicity can occur, namely that one factor could rotate to another factor. This possibility is illustrated in the following three examples.

Example 3.8.8. Let $d = 1$ and

$$f(x) = 1 - 2x^2 = (1 + \sqrt{2}x)(1 - \sqrt{2}x) = \varphi_1(x)\varphi_2(x).$$

Let $\sigma \in \mathrm{Gal}(\mathbb{Q}(\sqrt{2}) : \mathbb{Q})$ be given by $\sigma(\sqrt{2}) = -\sqrt{2}$. Then $\sigma(\varphi_1) = \varphi_2 = (-1) \cdot \varphi_1$. Now f_N is the product of $\zeta_N^j \cdot \varphi_k$ for $0 \leqslant j < N$ and $k = 1, 2$. If N is odd, then $-1 \notin \Omega_N$ and so all $2N$ factors are distinct. Our earlier arguments then show that $f_N = g_N$. However, if N is even, then $-1 \in \Omega_N$, and the set of rotates of φ_1 coincide with set of those of φ_2, and so $f_N = g_N^2$ for even N. Here f is an irreducible polynomial with a pair of roots whose ratio is a nontrivial root of unity.

The commingling of absolutely irreducible factors under rotations can happen in more subtle ways.

Example 3.8.9. Let $d = 1$ and $f(x) = 1 - 2x + 4x^2 - 3x^3 + x^4$, which is full and irreducible in $\mathbb{Z}[\mathbb{Z}]$. Let $\lambda = (1 + \sqrt{5})/2$, $\mu = (1 - \sqrt{5})/2$, and $\zeta = \zeta_5$. The absolutely irreducible factorization of f is

$$f(x) = (1 - \zeta\lambda x)((1 - \zeta^4\lambda x)(1 - \zeta^2\mu x)(1 - \zeta^3\mu x) = \varphi_1(x)\varphi_2(x)\varphi_3(x)\varphi_4(x).$$

Note that $\zeta^3 \cdot \varphi_1 = \varphi_2$ and that $\zeta \cdot \varphi_3 = \varphi_4$. If N is relatively prime to 5, then $\zeta \notin \Omega_N$, and so all $4N$ rotates are distinct and $f_N = g_N$ as before. However, if $5 \mid N$ then $\zeta \in \Omega_N$ and each rotate is repeated twice, and so $f_N = g_N^2$ in this case.

What is driving this example is the inclusion $\mathbb{Q}(\sqrt{5}) \subset \mathbb{Q}(\zeta)$, and so the Galois automorphism $\sqrt{5} \mapsto -\sqrt{5}$ of $\mathbb{Q}(\sqrt{5})$ is the restriction of the automorphism $\zeta \mapsto \zeta^2$ of $\mathbb{Q}(\zeta)$.

Remark 3.8.10. Irreducible polynomials in $\mathbb{Z}[x]$ having distinct roots whose ratio is a root of unity, such as those in the previous two examples, are called *degenerate*. Such polynomials have an extensive literature (see for instance [32, §1.1.9]), and appear in the celebrated Skolem-Mahler-Lech Theorem that the set of indices at which a recurring sequence of integers vanishes is, modulo a finite set, the union of arithmetic progressions [7].

There is a simple way to detect whether $f(x) \in \mathbb{Z}[x]$ is degenerate. Introduce a new variable t, and compute the resultant $g(x) \in \mathbb{Z}[x]$ of the polynomials $f(tx)$ and $f(t)$ with respect to t, which can be done efficiently using rational arithmetic. The roots of $g(x)$ are the ratios of all pairs of roots of f. Thus $f(x)$ is degenerate if and only if $g(x)$ contains a nontrivial cyclotomic factor. Applying this to $f(x)$ from the previous example gives

$$g(x) = (x-1)^5(x^4-4x^3+6x^2+x+1)(x^4+x^3+6x^2-4x+1)(x^4+x^3+x^2+x+1).$$

The last factor reveals that $f(x)$ has two roots whose ratio is a nontrivial 5th root of unity.

Example 3.8.11. Let $d = 2$ and $f(x,y) = 1-x-y-xy+x^2+y^2$, which is full and irreducible in $\mathbb{Z}[\mathbb{Z}^2]$. Let $\zeta = \zeta_3$. The absolutely irreducible factorization of f is

$$f(x,y) = (1+\zeta x+\zeta^2 y)(1+\zeta^2 x+\zeta y) = \varphi_1(x,y)\varphi_2(x,y).$$

Here φ_1 is mapped to φ_2 by the element σ in $\mathrm{Gal}(\mathbb{Q}(\zeta):\mathbb{Q})$ mapping ζ to ζ^2, and also $\sigma(\varphi_1) = \varphi_2 = (\zeta,\zeta^2)\cdot\varphi_1$. By the now familiar arguments, if N is relatively prime to 3 then $\zeta \notin \Omega_N$, and so all rotates are distinct and hence $f_N = g_N$. However, if 3 divides N, then distrinct rotates are repeated twice, and so $f_N = g_N^2$. For instance

$$f_3 = (1+3x^3+3y^3+3x^6-21x^3y^3+3y^3+x^9+3x^3y^6+3x^6y^3+y^9)^2 = g_3^2.$$

With these examples in mind, we come to the main result of this section.

Theorem 3.8.12. Let $f \in \mathbb{Z}[\mathbb{Z}^d]$ be irreducible, which we may assume is adjusted with positive constant term. For every $N \geqslant 1$ there is an irreducible $g_N \in \mathbb{Z}[N\mathbb{Z}^d]$ and $e_N \geqslant 1$ such that

$$\langle f\rangle_{\mathbb{Z}[\mathbb{Z}^d]} \cap \mathbb{Z}[N\mathbb{Z}^d] = \langle g\rangle_{\mathbb{Z}[N\mathbb{Z}^d]} \quad \text{and} \quad f_N = g_N^{e_N}.$$

The multiplicity e_N can be computed from the absolutely irreducible factorization of f in $\mathbb{C}[\mathbb{Z}^d]$. If the support of f generates a finite-index subgroup of $\mathbb{Z}^d$, then there is an integer $Q(f)$, which can also be computed from the absolutely irreducible factors of f, such that $e_N = 1$ for every N that is relatively prime to $Q(f)$. Finally,

$$\langle f^k\rangle_{\mathbb{Z}[\mathbb{Z}^d]} \cap \mathbb{Z}[N\mathbb{Z}^d] = \langle g_N^k\rangle_{\mathbb{Z}[N\mathbb{Z}^d]}$$

for every $k \geqslant 1$.

Proof. Recall our conventions in Remark 3.8.3. Let $\mathbb{K}$ be the splitting field of f, and $f = \widehat{f}(\mathbf{0})\varphi_1 \cdots \varphi_r$ be the factorization of f using monic absolutely irreducible $\varphi_j \in \mathbb{K}[\mathbb{Z}^d]$ from Proposition 3.7.3. Let $\Phi = \{\varphi_1, \ldots, \varphi_r\}$. Since the φ_j are monic, $\mathrm{Gal}(\mathbb{K} : \mathbb{Q})$ permutes the elements of Φ, and this action is transitive by irreducibility of f.

Now fix $N \geqslant 1$. Then $\mathbb{K}(\zeta_N)$ is a normal extension of $\mathbb{Q}$. Let $G = \mathrm{Gal}(\mathbb{K}(\zeta_N) : \mathbb{Q})$. Consider the set $\Omega_N^d \times \Phi$. The group Ω_N^d acts on this set via $\boldsymbol{\omega}' \cdot (\boldsymbol{\omega}, \varphi_j) = (\boldsymbol{\omega}'\boldsymbol{\omega}, \varphi_j)$. The group G also acts on this set via $\sigma \cdot (\boldsymbol{\omega}, \varphi_j) = (\sigma(\boldsymbol{\omega}), \sigma(\varphi_j))$. More precisely, $\sigma \in G$ acts of the first coordinate using its restriction to $\mathbb{Q}(\zeta_N)$ and on the second coordinate using its restriction to $\mathbb{K}$. These actions combine to give an action of the semidirect product $G \ltimes \Omega_N^d$ defined using the action of G on Ω_N^d, so that $\sigma\boldsymbol{\omega} = \sigma(\boldsymbol{\omega})\sigma$.

Define an equivalence relation $\sim$ on $\Omega_N^d \times \Phi$ by $(\boldsymbol{\omega}, \varphi_j) \sim (\boldsymbol{\omega}', \varphi_k)$ if and only if $\boldsymbol{\omega} \cdot \varphi_j = \boldsymbol{\omega}' \cdot \varphi_k$. It is routine to verify that $G \ltimes \Omega_N^d$ preserves equivalence classes. Since $\mathrm{Gal}(\mathbb{K} : \mathbb{Q})$ acts transitively on Φ, it follow that $G \ltimes \Omega_N^d$ acts transitively on $\Omega_N^d \times \Phi$. Hence all equivalence classes have the same cardinality, say $e_N \geqslant 1$. Pick one representative $(\boldsymbol{\omega}, \varphi_j)$ from each equivalence class, and let $\widetilde{g}_N$ be the product of the corresponding polynomials $\boldsymbol{\omega} \cdot \varphi_j$.

Observe that by its construction $\widetilde{g}_N$ is invariant under $G \ltimes \Omega_N^d$. Invariance under Ω_N^d implies that $\widetilde{g}_N \in \mathbb{K}(\zeta_N)[N\mathbb{Z}^d]$, and invariance under G further implies that $\widetilde{g}_N \in \mathbb{Q}[N\mathbb{Z}^d]$. Then transitivity of $G \ltimes \Omega_N^d$ on $\Omega_N^d \times \Phi$ shows that $\widetilde{g}_N$ is irreducible in $\mathbb{Q}[N\mathbb{Z}^d]$.

We have that $f_N = \widehat{f}(\mathbf{0})^{N^d}\, \widetilde{g}_N^{e_N}$. Let q be the least positive integer such that $q\widetilde{g}_N \in \mathbb{Z}[N\mathbb{Z}^d]$, so that $g_N := q\widetilde{g}_N$ is primitive. Then

$$f_N = \big(\widehat{f}(\mathbf{0})^{N^d}/q^{e_N}\big) g_N^{e_N}.$$

But both f_N and $g_N^{e_N}$ are primitive with positive constant terms, and hence $f_N = g_N^{e_N}$.

We now turn to computing e_N. Each of the absolutely irreducible factors φ_j has the same support since they are all Galois conjugates. Let Γ_φ denote the common support group of each. By Lemma 3.8.6, each contributes multiplicity $|\mathbb{Z}^d/(\Gamma_\varphi + N\mathbb{Z}^d)|$. Further multiplicity arises if one factor can be rotated by an element of Ω_N^d to another. This property divides Φ into equivalence classes, with all classes having the same cardinality s. It then follows that $e_N = |\mathbb{Z}^d/(\Gamma_\varphi + N\mathbb{Z}^d)|s$.

Next, we determine sufficient conditions on N so that $e_N = 1$. Assume that Γ_f has finite index in $\mathbb{Z}^d$. Clearly $\Gamma_f \subset \Gamma_\varphi$, and so Γ_φ also has finite index. By Lemma 3.8.6, if N is relatively prime to the index $[\mathbb{Z}^d : \Gamma_\varphi]$ of Γ_φ, then $|\mathbb{Z}^d/(\Gamma_\varphi + N\mathbb{Z}^d)| = 1$.

To analyze when one φ_j can rotate to another, we need to consider the group $\Omega_{\mathbb{K}}$ of roots of unity in the splitting field $\mathbb{K}$ of f. This is a finite cyclic group, and so equals Ω_n for some $n \geqslant 1$. Now $[\mathbb{Q}(\zeta_n) : \mathbb{Q}] = \varphi(n)$, where φ denotes the Euler function. Since $\mathbb{Q}(\zeta_n) \subset \mathbb{K}$, it follows that $\varphi(n) \leqslant [\mathbb{K} : \mathbb{Q}]$. A simple argument shows that $\varphi(n) \geqslant \sqrt{n}/2$ for all $n \geqslant 1$, and so $n \leqslant 4[\mathbb{K} : \mathbb{Q}]^2$. Hence if N is relatively prime to $(4[\mathbb{K} : \mathbb{Q}]^2)!$, then $\Omega_N \cap \Omega_{\mathbb{K}} = \{1\}$. For such an N suppose that $\boldsymbol{\omega} \cdot \varphi_i = \varphi_j$ for some $\boldsymbol{\omega} \in \Omega_N^d$. For each $\mathbf{k} \in \operatorname{supp} \varphi_i = \operatorname{supp} \varphi_j$ we have that $\boldsymbol{\omega}^{\mathbf{k}} \widehat{\varphi}_i(\mathbf{k}) = \widehat{\varphi}_j(\mathbf{k})$, and so

$$\boldsymbol{\omega}^{\mathbf{k}} = \widehat{\varphi}_j(\mathbf{k}) / \widehat{\varphi}_i(\mathbf{k}) \in \Omega_N \cap \Omega_{\mathbb{K}} = \{1\}.$$

But this implies that $\varphi_i = \varphi_j$.

Putting these together, we let $Q(f) = [\mathbb{Z}^d : \Gamma_\varphi](4[\mathbb{K} : \mathbb{Q}]^2)!$, and conclude that if N is relatively prime to $Q(f)$ then $e_N = 1$. □

3.9 Remarks and questions

Here we make some further remarks and ask several questions related to decimations.

3.9.1 More general lattices

Let us call a finite-index subgroup of $\mathbb{Z}^d$ a *lattice*. We have used the sequence $\{N\mathbb{Z}^d\}$ of lattices to define decimation, but these definitions easily extend to all lattices. Let $\Lambda \in \mathbb{Z}^d$ be a lattice, and let Ω_Λ denote the dual group of $\mathbb{Z}^d/\Lambda$, which has cardinality $[\mathbb{Z}^d : \Lambda]$, the index of Λ in $\mathbb{Z}^d$. Define $f_{\langle\Lambda\rangle} = \prod_{\boldsymbol{\omega} \in \Omega_\Lambda} \boldsymbol{\omega} \cdot f$, and

$$L_\Lambda f = E_{[\mathbb{Z}^d : \Lambda]}\Big(\frac{1}{[\mathbb{Z}^d : \Lambda]} \log |\widehat{f}_{\langle\Lambda\rangle}|\Big). \tag{3.17}$$

For a sequence $\{\Lambda_N\}$ of lattices, let us say $\Lambda_N \to \infty$ if for every $r > 0$ we have that $\{\mathbf{n} \in \Lambda_N : \|\mathbf{n}\| < r\} = \{\mathbf{0}\}$ for all large enough N.

Question 3.9.1. Let $0 \neq f \in \mathbb{C}[\mathbb{Z}^d]$, and let $\{\Lambda_N\}$ be a sequence of lattices with $\Lambda_N \to \infty$. Do the concave hulls $CH(L_{\Lambda_N})$ converge uniformly on $\mathcal{N}_f$ to D_f?

Our methods for $N\mathbb{Z}^d$ do not extend directly to this more general situation. We made essential use of the property of f_N that it is a polynomial in the Nth powers of the variables, enabling us to apply the Mahler estimates to the polynomial $E_N \widehat{f}_N$ of lower degree, gaining a crucial improvement. There is no corresponding argument for general lattices.

3.9.2 Partial decimation

By taking different sequences of lattices, we can in effect decimate along lower rank subgroups. The following example illustrates this idea.

Let $d = 2$ and $f(x, y) = 1 + x + y$. We will use the sequence of lattices $\Lambda_N = N\mathbb{Z} \oplus \mathbb{Z}$, which corresponds to decimating with respect to x. Using the notation from the previous section, $\Omega_{\Lambda_N} = \Omega_N \times \{1\}$, and so

$$f_{\langle \Lambda_N \rangle}(x, y) = \prod_{\omega \in \Omega_N} (1 + \omega x + y) = (1 + y)^N \pm x^N.$$

It is well-known that the growth rate of the binomial coefficients can be computed using Stirling's approximation to be

$$\frac{1}{N} \log \binom{N}{pN} \approx \eta(p) := -p \log p - (1 - p) \log(1 - p)$$

for $0 \leqslant p \leqslant 1$. Hence the decimation limit $D_f^{(1)}(r, s)$ of f with respect to x using this sequence of lattices is the concave hull of the curve $\big(0, p, \eta(p)\big)$ for $0 \leqslant p \leqslant 1$ together with the point $(1, 0, 0)$, as shown in Figure 3.7(a).

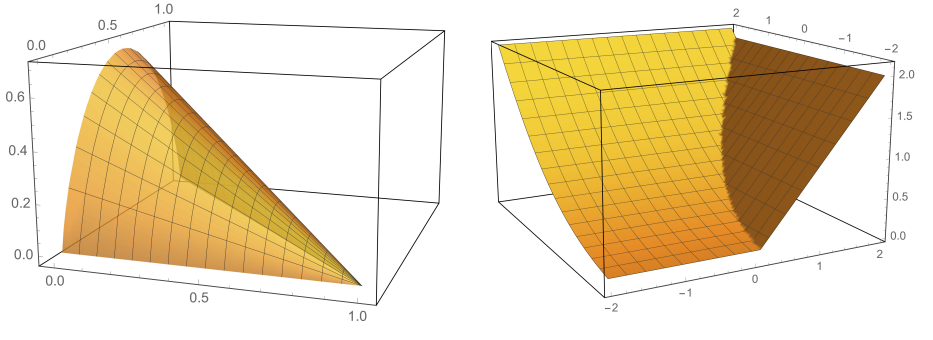

(a) (b)

Figure 3.7: (a) The partial decimation limit of $1 + x + y$, and (b) its partial Ronkin function

Define the partial Ronkin function of f with respect to x to be

$$R_f^{(1)}(u, v) = \int_0^1 \log |f(e^u e^{2\pi i\theta}, e^v)| \, d\theta.$$

Figure 3.7(b) shows this in our case. One can show that here $D_f^{(1)} = -\big(R_f^{(1)}\big)^*$ on $\mathcal{N}_f$.

This example suggests a more general phenomenon. Let $\mathcal{C}(\mathbb{Z}^d)$ denote the set of subgroups of $\mathbb{Z}^d$. We can give a topology to $\mathcal{C}(\mathbb{Z}^d)$ by declaring two subgroups to be close if they agree on a large ball around $\mathbf{0}$. For example, in this topology $\Lambda_N \to \{\mathbf{0}\}$ means $\Lambda_N \to \infty$ from §3.9.1, and in the above example $N\mathbb{Z} \oplus \mathbb{Z} \to 0 \oplus \mathbb{Z}$. This is a special case of the Chabauty topology on the set $\mathcal{C}(G)$ of closed subgroups of a locally compact group G. This topology is named after Claude Chabauty, who in 1950 introduced it [19] to generalize Mahler's compactness criterion [76] for lattices in $\mathbb{R}^d$ to lattices in locally compact groups. The Chabauty space $\mathcal{C}(G)$ has been investigated by many authors, for instance by Cornulier [21] when G is abelian. Even for familiar groups their Chabauty space can be intricate to analyze. For example, Hubbard and Pourezza [50] used a tricky argument to prove that $\mathcal{C}(\mathbb{R}^2)$ is homeomorphic to the four-dimensional sphere.

Let K be a compact subgroup of $\mathbb{T}^d$, and let μ_K denote normalized Haar measure on K. For $\mathbf{s} \in K$ we let $e^{2\pi i\mathbf{s}}\mathbf{u}$ mean $(e^{2\pi i s_1}u_1, \ldots, e^{2\pi i s_d}u_d)$. We then define the Ronkin function of f with respect to K to be

$$R_f^{(K)}(\mathbf{u}) = \int_K \log|f(e^{2\pi i\mathbf{s}}\mathbf{u})|\, d\mu_K(\mathbf{s}).$$

Question 3.9.2. Is there a limiting shape for decimations corresponding to a sequence of lattices $\{\Lambda_N\}$ in $\mathbb{Z}^d$ converging to a non-trivial subgroup $\Gamma \in \mathcal{C}(\mathbb{Z}^d)$?

3.9.3 Exponential size of decimation coefficients

In Example 3.2.3 we saw that if $f \in \mathbb{C}[\mathbb{Z}]$ is allowed to have complex coefficients, then some of the coefficients of f_N may have exponential size drastically different from that predicted by D_f. However, if $f \in \mathbb{Z}[\mathbb{Z}]$ is restricted to have integer coefficients, then this behavior cannot happen, as indicated by Example 3.2.2. More precisely, using the diophantine results of Gelfond mentioned there, one can show that if $f \in \mathbb{Z}[\mathbb{Z}]$ has $\operatorname{supp} f = \{0, 1, \ldots, r\}$ and $\varepsilon > 0$, then for all sufficiently large N we have that $|\widehat{f_N}(kN)|$ is between $e^{N(D_f(k)\pm\varepsilon)}$ for each $0 \leqslant k \leqslant r$ for which $\widehat{f_N}(kN) \neq 0$.

This raises the intriguing question of whether this extends to $f \in \mathbb{Z}[\mathbb{Z}^d]$ for $d \geqslant 2$, i.e., do all nonzero coefficients of f_N have the approximate exponential size predicted by D_f. The following gives a precise quantitative formulation.

Question 3.9.3. Let $f \in \mathbb{Z}[\mathbb{Z}^d]$. Fix $\mathbf{r}_0 \in \mathcal{N}_f$, and let $\varepsilon > 0$. Are there $\delta > 0$ and $N_0 \geqslant 1$ such that if $N \geqslant N_0$ and $\mathbf{r} \in N^{-d}\mathbb{Z}^d \cap \mathcal{N}_f$ with $\|\mathbf{r} - \mathbf{r}_0\| < \delta$,

and if $L_N f(\mathbf{r}) \neq -\infty$, then $|L_N f(\mathbf{r}) - D_f(\mathbf{r})| < \varepsilon$? Can δ and N_0 be chosen uniformly for $\mathbf{r}_0 \in \mathcal{N}_f$?

Some evidence for a positive answer comes from polynomials in two variables related to dimer models, as discussed in Remark 3.2.6. Using the additional machinery afforded by the physical interpretation of the related partition function and the resulting subadditivity, the exponential size of the coefficients can be shown to obey the estimates in the question. In particular, this applies to $f(x,y) = 1 + x + y$, although we do not know of any direct argument for this.

3.9.4 Continuity of $\exp[D_f]$ in the coefficients of f

Start by fixing a cube $B_n = \{-n, \dots, n\}^d \subset \mathbb{Z}^d$. We can identify a polynomial $f \in \mathbb{C}[\mathbb{Z}^d]$ whose support is in B_n with its coefficient function $\widehat{f} \in \mathbb{C}^{B_n}$. Boyd [11] showed that the function $\mathbb{C}^{B_n} \to [0,\infty)$ given by $\widehat{f} \mapsto M(f) = \exp[m(f)]$ is continuous in the coefficients of f.

Recalling that $m(f)$ is the maximum value of D_f, this suggests looking at $\exp[D_f]$, which is a nonnegative upper semicontinuous function on B_n (the discontinuities occur at the boundary of $\mathcal{N}_f \subset B_n$). A function $\varphi\colon B_n \to \mathbb{R}$ is upper semicontinuous if and only if its subgraph $\{(\mathbf{u}, t) \in B_n \times \mathbb{R} : t \leqslant \varphi(\mathbf{u})\}$ is closed in $B_n \times \mathbb{R}$. The space $\mathrm{USC}(B_n)$ of all upper semicontinuous functions on B_n carries a natural topology by declaring two elements to be close if their subgraphs are close in the Hausdorff metric on closed subsets of $B_n \times \mathbb{R}$ (see [6] for details).

Question 3.9.4. Is the map $\widehat{f} \to \exp[D_f]$ from $\mathbb{C}^{B_n}$ to $\mathrm{USC}(B_n)$ continuous?

3.9.5 Nonprincipal actions

Decimation makes sense for every algebraic $\mathbb{Z}^d$-action (indeed for every algebraic action of a countable residually finite group). Suppose that $\mathfrak{a}$ is an ideal in $\mathbb{Z}[\mathbb{Z}^d]$, and let $X_\mathfrak{a}$ be the dual group of $\mathbb{Z}[\mathbb{Z}^d]/\mathfrak{a}$ as described in §3.6. The commutative algebra there shows that the Nth decimation $r_N(X_\mathfrak{a})$ is defined by the contracted ideal $\mathfrak{a} \cap \mathbb{Z}[N\mathbb{Z}^d]$. However, there is no obvious replacement for g_N to measure growth when $\mathfrak{a}$ is not principal,

Question 3.9.5. If $\mathfrak{a}$ is a nonprincipal ideal in $\mathbb{Z}[\mathbb{Z}^d]$, are there objects related to the contractions $\mathfrak{a} \cap \mathbb{Z}[N\mathbb{Z}^d]$ which can be normalized to converge to a limiting object?

If $\mathfrak{a}$ is not principal, then the $\mathbb{Z}^d$-shift action on $X_\mathfrak{a}$ has zero entropy. However, by restricting the shift to iterates close to lower dimensional subspaces of $\mathbb{R}^d$ the action can have positive entropy [13, §6]. This suggests that the partial decimations from §3.9.2 may play a role here.

Examining concrete examples may shed some light on this question. These include the case of commuting toral automorphisms (see [52, §6] for many such examples), the $\mathbb{Z}^2$-action defined by multiplication by 2 and by 3 on $\mathbb{T}$ (corresponding to $\mathfrak{a} = \langle x-2, y-3\rangle$), and the so-called space helmet example [31, Example 5.8] (corresponding to $\mathfrak{a} = \langle 1 + x + y, z - 2\rangle$).

An important example of a different character is due to Ledrappier [61], which corresponds to the nonprincipal ideal $\langle 1 + x + y, 2\rangle \subset \mathbb{Z}[\mathbb{Z}^2]$. This example has zero entropy as a $\mathbb{Z}^2$-action, but strictly positive entropy along every 1-dimensional subspace of $\mathbb{R}^2$ (see [13, Example 6.4] for the explicit description). Another curious feature of this example is decimation self-similarity. Because $(1+x+y)^{2^n} = 1+x^{2^n}+y^{2^n}$ when taken mod 2, the 2^nth decimation of the example, when rescaled by 2^n, is just the original action.

3.10 Example of computing the decimation limit

There are few explicit calculations of the logarithmic Mahler measure, or more generally of the Ronkin function, of polynomials in $\mathbb{Z}[\mathbb{Z}^d]$ when $d \geqslant 2$. Depending on the relative sizes of the coefficients, evaluation of the integrals involved typically requires the torus to be subdivided into a large number of subregions with complicated boundaries, and so simple formulas in terms of familiar functions are rare.

Here we treat the case $f(x, y) = 1 + x + y$ from Example 3.2.4, where these calculations can be carried out, resulting in the formulas (3.8) and (3.9) for D_f.

Smyth [105] first computed the logarithmic Mahler measure $m(f) = R_f(0, 0)$ to have the value in (3.10). Twenty years later Maillot [78, §7.3], aided by Cassigne, computed the entire Ronkin function $R_f(u, v)$, providing in his long memoir a concrete example of the canonical height of a hypersurface. Their result involves the Bloch-Wigner dilogarithm function, which is an alternative formulation of the series representation in our formulas. Lundqvist [70] gave the formulas for the partial derivatives of R_f we use here. He also investigated the polynomial $1 + x + y + z$, and showed that the second order partial derivatives of its Ronkin function can be expressed in terms of standard elliptic functions.

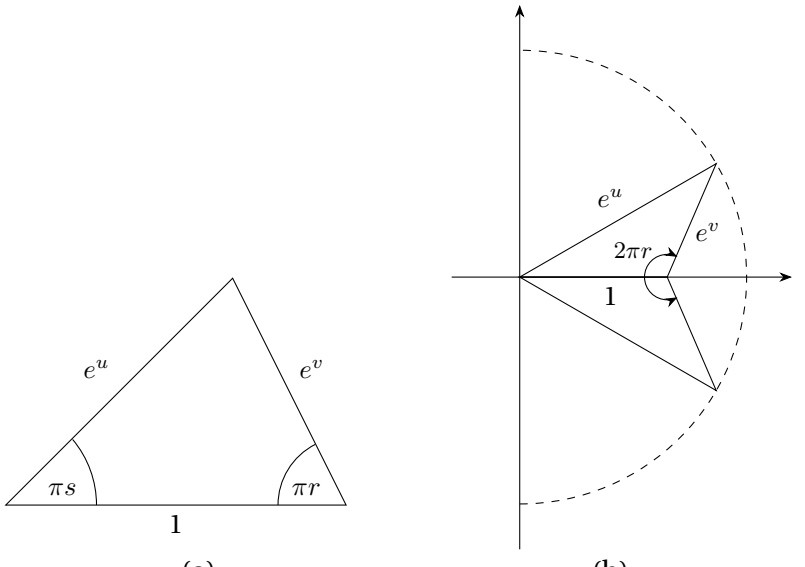

Figure 3.8: Determining partial derivatives from angles and sides

Let $\Delta = \mathcal{N}_f$ be the unit simplex, and denote its interior by Δ°. Let $\mathcal{A}_f$ be the amoeba of f, as shown in Figure 3.5, and $\mathcal{A}_f^\circ$ be its interior. To evaluate $R_f^*(r,s)$ for $(r,s) \in \Delta^\circ$, we need to know the value of $(u,v) \in \mathcal{A}_f^\circ$ at which the partial derivatives of $R_f(u,v)$ with respect to u and v equal r and s, respectively. Fortunately, there is a simple relationship that was established by Lundqvist [70], whose treatment we follow.

Lemma 3.10.1. Let $(u,v) \in \mathcal{A}_f^\circ$, so that 1, e^u, and e^v form the sides of a nondegenerate triangle. Let πr and πs be the angles in this triangle shown in Figure 3.8(a). Then

$$\frac{\partial R_f}{\partial u}(u,v) = r \quad \text{and} \quad \frac{\partial R_f}{\partial v}(u,v) = s. \tag{3.18}$$

Proof. We will compute the partial derivatives by differentiating the integrand in

$$\begin{aligned} R_f(u,v) &= \int_0^1 \int_0^1 \log|1 + e^u e^{2\pi i\theta} + e^v e^{2\pi i\varphi}| \, d\theta \, d\varphi \\ &= \mathrm{Re}\Big[\int_0^1 \int_0^1 \log(1 + e^u e^{2\pi i\theta} + e^v e^{2\pi i\varphi}) \, d\theta \, d\varphi\Big]. \end{aligned}$$

In the last line $\log$ represents a local inverse to $\exp$, which is well-defined up to the addition of an integral multiple of $2\pi i$. After taking partial derivatives, we will get a result that is independent of this multiple.

By symmetry, it suffices to compute $\partial R_f/\partial u$. Differentiating the integrand gives

$$\frac{\partial R_f}{\partial u}(u,v) = \mathrm{Re}\Big[\int_0^1\int_0^1 \frac{e^u e^{2\pi i\theta}}{1+e^u e^{2\pi i\theta}+e^v e^{2\pi i\varphi}}\,d\theta\,d\varphi\Big].$$

Rewriting the integrals as contour integrals, we see that

$$\begin{aligned}\int_0^1\int_0^1 \frac{e^u e^{2\pi i\theta}}{1+e^u e^{2\pi i\theta}+e^v e^{2\pi i\varphi}}\,d\theta\,d\varphi &= \frac{1}{(2\pi i)^2}\int_{|z|=e^u}\int_{|w|=e^v}\frac{1}{1+z+w}\,dz\frac{dw}{w}\\ &= \frac{1}{2\pi i}\int_{|w|=e^v}\Big[\frac{1}{2\pi i}\int_{|z|=e^u}\frac{dz}{z-(-1-w)}\Big]\frac{dw}{w}.\end{aligned}$$

The inner integral is the winding number of the circle of radius e^u around $-1-w=-1-e^v e^{2\pi i\varphi}$, and so has value 1 if $|1+e^v e^{2\pi i\varphi}|<e^u$ and 0 if $|1+e^v e^{2\pi i\varphi}|>e^u$ (these are mistakenly reversed in [70]). A glance at Figure 3.8(b) shows that the value is 1 for an interval of φ of length $2\pi r$, and 0 otherwise. Since $(1/2\pi i)(dw/w)$ is normalized Lebesgue measure $d\varphi$, we obtain that $(\partial R_f/\partial u)(u,v)=r$. □

To compute the decimation limit D_f, we need to express u and v in terms of r and s. Let $a=e^u$ and $b=e^v$ be the sides of the triangle in Figure 3.8(a). By the law of sines,

$$\frac{a}{\sin\pi r}=\frac{b}{\sin\pi s}=\frac{1}{\sin\pi(1-r-s)}=\frac{1}{\sin\pi(r+s)},$$

and hence

$$a=a(r,s)=e^{u(r,s)}=\frac{\sin\pi r}{\sin\pi(r+s)}, \tag{3.19}$$

$$b=b(r,s)=e^{v(r,s)}=\frac{\sin\pi s}{\sin\pi(r+s)}. \tag{3.20}$$

For $(u,v)\in\mathcal{A}_f^\circ$ it follows from the definition (3.11) that

$$-R_f^*(u,v)=\inf_{(r,s)\in\Delta^\circ} R_f(u,v)-ru-sv,$$

and by calculus the infimum is attained at the (u,v) given by (3.18). Thus for $(r,s)\in\Delta^\circ$ we have that

$$D_f(r,s)=-R_f^*\big(u(r,s),v(r,s)\big)=R_f\big(u(r,s),v(r.s)\big)-r\,u(r,s)-s\,v(r,s), \tag{3.21}$$

where $u(r,s)$ and $v(r,s)$ are determined by (3.19) and (3.20).

Remark 3.10.2. Observe that the functions $u(r,s)$ and $v(r,s)$ in (3.19) and (3.20) are real analytic on Δ°. Also, $R_f(u,v)$ is real analytic on $\mathcal{A}_f^\circ$. Together these show that $D_f(r,s)$ is real analytic on Δ°.

It remains to compute $R_f(u,v)$. By symmetry it suffices to assume that $u \geqslant v$. Using Jensen's formula (3.5), we see that

$$\begin{aligned} R_f(u,v) &= \int_0^1 \int_0^1 \log|1 + e^u e^{2\pi i\theta} + e^v e^{2\pi i\varphi}| \, d\theta \, d\varphi \\ &= u + \int_0^1 \int_0^1 \log|e^{-u} + e^{v-u} e^{2\pi i\varphi} + e^{2\pi i\theta}| \, d\theta \, d\varphi \\ &= u + \int_0^1 \log^+|e^{-u} + e^{v-u} e^{2\pi i\varphi}| \, d\varphi. \end{aligned}$$

Note that $|e^{-u} + e^{v-u}e^{2\pi i\varphi}| \geqslant 1$ if and only if $|1 + e^v e^{2\pi i\varphi}| \geqslant e^u$, and another glance at Figure 3.8(b) shows this occurs exactly when $-\pi(1-r) \leqslant 2\pi\varphi \leqslant \pi(1-r)$. Hence

$$\begin{aligned} R_f(u,v) &= u + \int_{-\frac{1}{2}(1-r)}^{\frac{1}{2}(1-r)} \log|e^{-u} + e^{v-u} e^{2\pi i\varphi}| \, d\varphi \\ &= u - (1-r)u + \int_{-\frac{1}{2}(1-r)}^{\frac{1}{2}(1-r)} \log|1 + e^v e^{2\pi i\varphi}| \, d\varphi \\ &= r\,u + \int_{-\frac{1}{2}(1-r)}^{\frac{1}{2}(1-r)} \log|1 + e^v e^{2\pi i\varphi}| \, d\varphi. \end{aligned}$$

First suppose that $e^v < 1$, which corresponds to $(r,s) \in \Delta_1^\circ$, where Δ_1 is defined in (3.6). The series expansion of $\log(1+z)$ for $1+z$ in the domain of integration converges uniformly, and the imaginary part vanishes by symmetry. Hence

$$\begin{aligned} R_f(u,v) &= r\,u + \int_{-\frac{1}{2}(1-r)}^{\frac{1}{2}(1-r)} \sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{n} e^{nv} e^{2\pi i n\varphi} \, d\varphi \\ &= r\,u + \sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{n} e^{nv} \frac{1}{\pi n} \sin[\pi n(1-r)]. \end{aligned}$$

Recalling that $e^{v(r,s)} = b(r,s) = (\sin \pi s)/\sin[\pi(r+s)]$, we conclude that

$$\begin{aligned} D_f(r,s) &= R_f\big(u(r,s), v(r,s)\big) - r\,u(r,s) - s\,v(r,s) \\ &= \sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{\pi n^2} b(r,s)^n \sin[\pi n(1-r)] - s\,\log[b(r,s)]. \end{aligned} \tag{3.22}$$

Now suppose that $e^v > 1$, which corresponds to $(r,s) \in \Delta_2^\circ$, where Δ_2 is defined by (3.7). Then $\log|1 + e^v e^{2\pi i\varphi}| = v + \log|1 + e^{-v}e^{-2\pi i\varphi}|$. Calculating as before,

$$\begin{aligned} R_f(u,v) &= r\,u + (1-r)v + \int_{-\frac{1}{2}(1-r)}^{\frac{1}{2}(1-r)} \sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{n} e^{-nv} e^{-2\pi i n\varphi}\, d\varphi \\ &= r\,u + (1-r)v + \sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{\pi n^2} b(r,s)^{-n} \sin[\pi n(1-r)]. \end{aligned}$$

Thus for $(r,s) \in \Delta_2^\circ$ we find that

$$D_f(r,s) = \sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{\pi n^2} b(r,s)^{-n} \sin[\pi n(1-r)] + (1-r-s)\log[b(r,s)]. \quad (3.23)$$

Finally, note that on the overlap $\Delta_1 \cap \Delta_2$ inside Δ°, we have that $b(r,s) = 1$ and so the series in (3.6) and (3.7) converge and agree, hence give the value of $D_f(r,s)$ by continuity of the Legendre transform.