Convolutions of inverse linear functions via multivariate residues

Yuliy Baryshnikov [1]
Robin Pemantle [2,3]

ABSTRACT:
Let $F(z_1, \ldots z_d) = \prod_{j=1}^n \frac{\eta}{l_j(z_1, \ldots, z_d)^{n_j}}$ be the quotient of an analytic function by a product of linear functions $l_j := 1 - \sum b_{ij} z_i$. We compute asymptotic formulae for the Taylor coefficients of F via the multivariate residue approach begun by [BM93]. By means of stratified Morse theory, we are able to give a short and fully implementable algorithm for determining an asymptotic series expansion.

Keywords: queuing, asymptotic, generating function, Morse theory.

Subject classification: Primary: 32A27, 60K25, 14Q20. secondary: 14Q20, 90B22.

[1] Bell Laboratories, Lucent Technologies, 700 Mountain Avenue, Murray Hill, NJ 07974-0636, ymb@research.bell-labs.com

[2] Research supported in part by National Science Foundation grant # DMS 0103635

[3] University of Pennsylvania, Department of Mathematics, 209 S. 33rd Street, Philadelphia, PA 19104, pemantle@math.upenn.edu

1 Introduction

Let $\mathbf{z}$ denote the vector of complex numbers $(z_1, \ldots, z_d)$ with standard multi-index notation $\mathbf{z}^\mathbf{r} := z_1^{r_1} \cdots z_d^{r_d}$. Let

$$F(\mathbf{z}) = \sum a_\mathbf{r} \mathbf{z}^\mathbf{r} = \frac{\eta(\mathbf{z})}{\prod_{j=1}^n l_j(\mathbf{z})^{n_j}} \tag{1.1}$$

be a d-variable generating function which is the quotient of an analytic function by a product of linear terms $l_j = 1 - \sum_{i=1}^d b_{ij} z_i$, with b_{ij} being any real numbers. These arise, among other places, in enumeration problems (see [DS03]), in queuing theory (see [BM93] and [KY96]) and in Markov modeling (see [Kar02]).

Pemantle and Wilson [PW04] give a method for asymptotic analysis of the coefficients $a_\mathbf{r}$ in the general setting where F is a rational function with a pole variety that is locally completely reducible to unions of smooth sheets. Their methods compute the asymptotics in many cases of interest, but have two shortcomings. First, their hypotheses require that the singularity of F which "controls" the asymptotics be a minimal singularity, that is, it must lie on the boundary of the domain of convergence of F. This hypothesis is considerably weaker than the hypotheses in [BR83] on behavior of F at the radius of convergence and nonnegativity of the coefficients $a_\mathbf{r}$. Nevertheless, the hypothesis sometimes fails (see Example 4 below, taken from the dimensional analysis of the stationary measure for a Markov model of TCP-IP protocol). Secondly, the methods of [PW04] are not computationally effective. Their method is to write the coefficients as multivariate complex integrals, to isolate the region contributing the most to the integral, and to evaluate the local integral by means of several integrating tricks. This relies on first identifying the correct minimal point where the denominator of F vanishes, for which they have a geometric characterization which is neither universal nor effective.

Our purpose in this paper is to give a complete and effective asymptotic analysis of the coefficients $a_\mathbf{r}$ in (1.1). Our methods combine the work of [PW04] with the approach outlined by Bertozzi and McKenna in [BM93]. They point out that the integral for $a_\mathbf{r}$ may be represented as a sum of integrals over basic homology cycles, among which some linear relations hold. They find enough relations to evaluate the integral in a few low-dimensional cases, and suggest that a more systematic treatment might prove worthwhile. Indeed, we find an algorithm that gives a full asymptotic series for $a_\mathbf{r}$ without the nonnegativity assumption. Asymptotics are as $\mathbf{r} \to \infty$ and are uniform as long as $\mathbf{r}/|\mathbf{r}|$ remains in a compact set whose complement has codimension 1. For asymptotics in this smaller set of "non-generic" directions, our methods do not apply and the methods of [PW04] are required.

We do not pick up directly where Bertozzi and McKenna left off, which would be to complete the analysis of the linear relations among the naturally occurring d-dimensional homology cycles. Instead, we use stratified Morse theory to compute directly the d-cycles that determine the asymptotic series for $a_\mathbf{r}$. A brief description of this program is as follows.

1. Write $a_{\mathbf{r}}$ as a Cauchy integral

$$a_{\mathbf{r}} = \left(\frac{1}{2\pi i}\right)^d \int_T \mathbf{z}^{-\mathbf{r}} F(\mathbf{z}) \frac{d\mathbf{z}}{\mathbf{z}} \tag{1.2}$$

where the torus T is a product of sufficiently small circles around the origin in each coordinate.

2. The torus T may be replaced by a homologous d-cycle in $(\mathbb{C}^*)^d$ minus the poles of F. Specifically, we denote by $-\infty$ the set where the integrand in (1.2) is sufficiently small, and represent T in the homology of $(\mathbb{C}^*)^d$ minus the poles of F, relative to $-\infty$.

3. Stratified Morse theory identifies these homology classes with saddles of the gradient $\mathbf{r} \log \mathbf{z}$ of the function $\mathbf{z}^{\mathbf{r}}$. Each such saddle lives in a stratum of dimension $j < d$ and yields a contribution which is an integral over a product of a cycle $\mathsf{cyc}^{\|}$ in the stratum with a cycle $\mathsf{cyc}^{\perp}$ in a transversal to the stratum.

4. A nonzero contribution at a saddle σ occurs when the vector $\mathbf{r}$ is in a certain positive cone determined by the geometry of the arrangement near σ.

5. The integral over $\mathsf{cyc}^{\perp}$ is equal to an easily computed spline, and the integral over $\mathsf{cyc}^{\|}$ is then asymptotically evaluated by the saddle point method.

Step 1 is common to [PW01, PW04, BM93, KY96]. This step is valid for completely general F, not just for functions with linear poles, as considered in this paper. Step 2 is done in [BM93]. Step 3 is new here. Both of these steps are valid again in complete generality, but the description of the stratification is significantly easier in the linear case. Step 4, is also new, at least in this level of generality. We have some idea of how this might be computed in the nonlinear case, but this is work in progress (see [HP02]). The transverse part of Step 5 is found in the literature in various forms, and this instance draws on [PW04] and [Pem00]. The description of the cones in Step 4 and the transverse integrals arising there, in the special case of a central arrangement with common intersection a single point, may be found in [BV99]; in fact some of the lemmas appearing in Section 5 were proved in another form in [BV99]. The saddle point part is common to all methods.

The remainder of the paper is organized as follows. In the next section we set up definitions and notation and state our main results, which are asymptotic formulae for $a_{\mathbf{r}}$. Section 3 gives a more detailed development of notation, explaining some of the terms in the statements of results and providing examples. Section 4 goes over some topological facts needed for the evaluation of the multidimensional complex contour integrals. Section 5 proves the main results. Section 6 then discusses how, algorithmically, to compute these formulae. The last section is concerned with extensions and generalizations. We end this introductory section by re-casting the steps 1 to 5 above as a sequence of calls to subroutines specified in Section 6.

Algorithm 1.1 (Main routine)

1. *Compute the intersection lattice of the flats defined by* l_j
2. *Find the unique saddle point in each flat*
3. *Sort the saddles by the height function* $f_{\mathbf{r}}$
4. *Set* σ *equal to the highest saddle, and until the leading saddle is identified, do:*
   (a) *Compute the positive cone* K_σ *at* σ*;*
   (b) *If* $\mathbf{r} \notin K_\sigma$ *then set* σ *to the next lower saddle and repeat do loop;*
   (c) *compute spline function* p_σ *at* σ*;*
   (d) *compute the polynomial specialization for the chamber containing* $\mathbf{r}$*;*
   (e) *if this is identically zero then set* σ *to the next lower saddle and repeat the loop (this happens when* η *is in a certain ideal defined later: in the case of a generic intersection of hyperlanes containing* σ *the condition is simply* $\eta(\sigma) = 0$*);*
   (f) *if* $\mathbf{r}$ *is on the boundary of a chamber then output that* $\mathbf{r}$ *is non-generic and halt;*
   (g) *identify* σ *as the leading saddle;*
5. *Compute stationary phase expansion by integrating coefficients of the polynomial.*

2 Statements of results

Let F be given by (1.1) in terms of l_j, $1 \leq j \leq n$, which are in turn defined by constants b_{ij}. Let planes denote the union of the coordinate planes and $\mathsf{poles} := \{\mathbf{z} : \prod_{j=1}^n l_j(\mathbf{z}) = 0\}$ denote the set of poles of F. Let $\xi := \mathbf{z}^{-\mathbf{r}-\mathbf{1}} F(\mathbf{z})\, d\mathbf{z}$ denote the integrand in (1.2).

Stratified Morse Theory tells us that the surface of integration, T, in equation (1.2) may be pushed down (with respect to the gradient field of $\mathbf{z}^{-\mathbf{r}}$) until it snags on a critical point of some stratum. We denote by Σ the set of all critical points of strata and let $S(\sigma)$ denote the stratum for which σ is a critical point (well defined for generic $\mathbf{r}$). The motivation for the definitions in the next section is that we are trying to find a cycle homologous to T over which the integral in (1.2) may be evaluated. It turns out there is a set $\Sigma(S)$ of saddles associated with each flat S of the hyperplane arrangement; letting $\Sigma := \bigcup_{S \in \mathcal{A}} \Sigma(S)$, there is a subset $\mathsf{contrib} \subseteq \Sigma$ such that the sum of certain quasi-local cycles $\mathsf{cyc}(\sigma)$ over $\mathsf{contrib}$ is homologous to T (a quasi-local cycle is an actual d-cycle which has certain locality properties).

The purpose of this section is to state the main results that go into the justification of Algorithm 1.1, with as little overhead as possible. There are a number of terms, whose full definitions are given in a later section, that we define here only by suggestive names. The main purpose of Section 3 will be to give precise definitions for terms other than those in the following definitions:

Definition 1 (height function) *Given a vector* $\mathbf{r}$ *of positive reals, define the height function* $f = f_{\mathbf{r}}$ *by*

$$f_{\mathbf{r}}(\mathbf{x}) := |\mathbf{r}| f_{\hat{\mathbf{r}}}(\mathbf{x}) = -\sum_{j=1}^{d} r_j \log |x_j| \, . \tag{2.1}$$

The saddles are critical points for this function, hence their location depends on $\mathbf{r}$. The height function is strictly convex in each orthant of $\mathbb{R}^d$.

Definition 2 (manifold of holomorphy of ξ) *Let* planes *denote the union of the coordinate hyperplanes and* poles *denote the union of hyperplanes in* $\mathcal{A}$*. Set* $\mathcal{M} := \mathbb{C}^d \setminus (\mathsf{planes} \cup \mathsf{poles})$ *to be the domain of holomorphy of* ξ*. We let* $\mathcal{M}_c := \{\mathbf{x} \in \mathcal{M} : f_{\hat{\mathbf{r}}}(\mathbf{x}) \leq c\}$*. The symbol* $H_*(\mathcal{M}, -\infty)$ *is used to denote the homology of the one-point compactification of* $\mathcal{M}$ *relative to infinity, or equivalently, to the inverse limit of* $(\mathcal{M}, \mathcal{M}_c)$ *as* $c \to -\infty$*. We will see that these spaces are homotopy equivalent once* $c < \mathsf{low}$*, the least critical value of* $f_{\hat{\mathbf{r}}}$*.*

In the next section, we will define the cone $\mathsf{gen} \subseteq (\mathbb{R}^+)^d$ of generic values for the $\mathbf{r}$ vector, whose complement non has codimension 1. All asymptotic results in this paper assume $\mathbf{r} \in \mathsf{gen}$ and enjoy the following uniformity:

Asymptotics are uniform as $\mathbf{r} \to \infty$ with $\mathbf{r}/|\mathbf{r}|$ remaining in a compact subset of gen. (2.2)

The poles of the function F form a hyperplane arrangement, the combinatorics and geometry of which largely determine the asymptotic behavior of the coefficients. The first two theorems, Theorems 2.1 and 2.2, involve a simplifying assumption of genericity of this arrangement, which is removed in the next two theorems. Specifically, we say that a hyperplane arrangement is generic if $j \leq d$, any set of planes of cardinality j from among the planes of the arrangement intersects in a flat of dimension $d - j$, and no set of more than d of these planes has a common intersection. All examples in [BM93] satisfy the genericity assumption, the most complicated one of which [BM93, Example 4.6] has $n = d = 3$.

Our first result gives a representation of the class $[T]$ in terms of homology classes associated with each saddle. The set $\mathsf{contrib}$ is defined in Definition 4 below for all arrangements and all generic $\mathbf{r}$. The cycle $\mathsf{cyc}^{\parallel}(\sigma)$ is defined in Definition 3 for all arrangements as well, while the cycle $\mathsf{cyc}^{\perp}(\sigma)$ is defined there only for generic arrangements.

Theorem 2.1 (decomposition of T) *Suppose that the poles of* F *form a generic hyperplane arrangement. Then for* $\mathbf{r} \in \mathsf{gen}$*,*

$$[T] = \sum_{\sigma \in \mathsf{contrib}} \mathbf{sgn}(\sigma) \, \mathsf{cyc}(\sigma) \, . \tag{2.3}$$

in $H_d(\mathcal{M}, -\infty)$. Here $\mathbf{sgn}(\mathbf{x})$ *is the sign of the product of the coordinates of* $\mathbf{x}$ *and* $\mathsf{cyc}(\sigma) = \mathsf{cyc}^{\perp}(\sigma) \times \mathsf{cyc}^{\parallel}(\sigma) \in \mathcal{S}^{\perp}(\sigma) \times \mathcal{S}^{\parallel}(\sigma)$.

The next result shows how to integrate: first in the normal direction, which is an iterated residue, then in the tangential direction. Suppose U and V are orthogonally complementary subspaces of $\mathbb{R}^d$, with orthonormal bases $\{u_1, \ldots, u_k\}$ and $\{v_1, \ldots, v_{d-k}\}$. Embed $\mathbb{R}^d$ in $\mathbb{C}^d$ and let $u'_j = iu_j$ and $v'_j = iv_j$. Let dx_j denote the complexified 1-form $du_j + idu'_j$ and $dy_j = dv_j + idv'_j$. Then $d\mathbf{z}^{\perp} := dx_1 \wedge \cdots \wedge dx_k$ and $d\mathbf{z}^{\parallel} := dy_1 \wedge \cdots \wedge dy_{d-k}$ are holomorphic forms and $d\mathbf{z}^{\perp} \wedge d\mathbf{z}^{\parallel} = \pm d\mathbf{z}$, where $d\mathbf{z}$ is the standard holomorphic volume form. The projection functions onto U and V respectively are denote $\mathbf{z}^{\perp}$ and $\mathbf{z}^{\parallel}$.

Theorem 2.2 (evaluation of the normal integral) *Suppose that* $\mathbf{r} \in \mathsf{gen}$*, that the arrangement of poles is generic, and that* $n_j \equiv 1$*. For* $\sigma \in \mathsf{contrib}$*, let*

$$d\mathbf{z} = d\mathbf{z}^{\perp} \wedge d\mathbf{z}^{\parallel}$$

denote the decomposition of the volume form into components parallel and orthogonal to the flat V_S *associated with* $S(\sigma)$*. Then*

$$\int_{\mathsf{cyc}(\sigma)} \xi \quad = \quad (2\pi i)^{\mathrm{codim}\,\sigma} \det{}_{\sigma}(\mathbf{b})^{-1} \int_{\mathsf{cyc}^{\parallel}(\sigma)} \mathbf{z}^{-\mathbf{r}-\mathbf{1}} \frac{\eta(\mathbf{z})}{\prod_{j \notin S(\sigma)} l_j(\mathbf{z})} \, d\mathbf{z}^{\parallel}$$

At this point, the asymptotic evaluation of $a_{\mathbf{r}}$ is reduced to a standard saddle point integral (see Lemma 5.5 in Section 5.3).

A representation of the homology class of T is important only up to its evaluation at the cohomology element corresponding to the integrand, ξ. We give a less constructive version of Theorem 2.1, valid when the hyperplane arrangement defined by the l_j is not generic, though Theorem 2.4 maintains the same computing power as Theorem 2.2.

Theorem 2.3 (decomposition of T, non-generic case) *Suppose that* $\mathbf{r} \in \mathsf{gen}$*. With no assumption on* $\mathcal{A}$*, Definition 3 may be extended to define cycles* $\mathsf{cyc}^{\perp}(\sigma)$ *such that*

$$[T] = \sum_{\sigma \in \mathsf{contrib}} \mathbf{sgn}(\sigma) \, \mathsf{cyc}(\sigma)$$

with $\mathsf{cyc}(\sigma) = \mathsf{cyc}^{\perp}(\sigma) \times \mathsf{cyc}^{\parallel}(\sigma)$*. If* $\sigma \in \mathsf{contrib}$ *then* $\mathsf{cyc}(\sigma)$ *is not homologous to zero in* $H_d(\mathcal{M}, -\infty)$*.*

Define a grading function on saddles by

$$|\sigma| = \sum_{j : \sigma \in V_j} n_j \, .$$

Theorem 2.4 (normal integral, non-generic case) *For any stratum $\mathcal{S} = \mathcal{S}_S$, one can define a function $P_\mathcal{S} : \mathcal{S} \times \mathcal{S}^\perp \to \mathbb{C}$ (see Definition 6 in Section 3) having the following properties.*

1. *P is piece-wise analytic and in fact piece-wise polynomial on any fiber $\mathbf{x} \times \mathcal{S}^\perp \to \mathbb{C}$, with support in the cone $K_S(\mathbf{x}) \subset \mathbf{x} \times \mathcal{S}^\perp$.*

2. *The degree of P_σ in $\mathbf{r}$ satisfies*
$$\deg P_\sigma \leq |\sigma| - \operatorname{codim} \sigma \tag{2.4}$$
and the homogeneous part of degree $|\sigma| - \operatorname{codim}\sigma$ is given by
$$\eta(\mathbf{x}) \left(\prod_{j \notin S(\sigma)} l_j(\mathbf{x}) \right)^{-n_j} p_\sigma$$
with p_σ defined in Definition 5.

3. $$\int_{\mathsf{cyc}(\sigma)} \xi \;=\; \mathbf{sgn}(\sigma) \int_{\mathsf{cyc}^\parallel(\sigma)} P_{\sigma, \mathbf{z}^\parallel} \, d\mathbf{z}^\parallel \tag{2.5}$$
if $\mathbf{r}$ is in the set $\mathsf{contrib}$ of Theorem 2.2, and is zero otherwise.

Our main result is that Algorithm 1.1 correctly yields asymptotic formulae for $a_\mathbf{r}$. This rests on the following theoretical result, which is essentially the union of Theorems 2.3, 2.4 and a standard result on Saddle point integration, Lemma 5.5.

Theorem 2.5 (main result) *Let F, $\mathcal{A}$, V_S, $f_\mathbf{r}$ and σ be as above. Examine the saddles in decreasing order of the values $f_\mathbf{r}$. Let σ^* be the first one in $\mathsf{contrib}$ for which $P_{\sigma,\mathbf{x}}$ is not identically zero, the condition for which is that η not lie in the ideal $\mathcal{I}(\sigma)$.*

Then $a_\mathbf{r}$ has a computable full asymptotic expansion, uniform over compact subcones of gen (cf. 2.2), whose leading behavior, in the case $\eta(\sigma^*) \neq 0$, is given by
$$a_\mathbf{r} = (2\pi)^{-\frac{1}{2}\dim \sigma^*} (1 + o(|\mathbf{r}|^{-1})) \frac{\eta(\sigma^*)}{\prod_{j : l_j(\sigma^*) \neq 0} l_j(\sigma^*)} p_{\sigma^*}(\mathbf{r}) \Lambda_{\sigma^*}(\mathbf{r})^{-1/2} (\sigma^*)^{-\mathbf{r}} \,. \tag{2.6}$$
If there is more than one point σ^* which maximizes $f_\mathbf{r}$ and for which $P_{\sigma^*,\cdot}$ not identically zero, then we must sum (2.6) over all such σ^*. In this case periodic cancellation is possible, so the sum of the leading terms may not always be the leading term of the sum.

3 Definitions and examples

We begin the description of notation with a table summarizing the quantities to be defined and listing which quantities depend on $\mathbf{r}$, which depend only on the unit vector $\hat{\mathbf{r}}$ in the direction $\mathbf{r}$, which depend on the saddle σ or only on its flat, S, and which depend only on the given data, F.

Notation	**Name**	**Depends on**
$\mathbf{b}$	matrix of coefficients	
$\det_\sigma(\mathbf{b})$	restricted determinant	$S(\sigma)$
$\mathbf{b}_j, \tilde{\mathbf{b}}_j$	the j^{th} normal and logarithmic normal	$\mathbf{x}$
$\mathcal{A}$	arrangement lattice	
V_S	the flat associated with $S \in \mathcal{A}$	S
$\mathcal{S}_S$	the stratum associated with S	S
Σ	the set of all saddles	$\mathbf{r}$
$S(\sigma)$	the S for which $\sigma \in \mathcal{S}_S$	$\hat{\mathbf{r}}$, σ
$\dim \sigma$	dimension of the flat $S(\sigma)$	S
$\mathcal{S}^\perp$	the normal to $\mathcal{S}_S$	S
$\mathsf{high}, \mathsf{low}$	max and min critical values	$\hat{\mathbf{r}}$
$\mathsf{cyc}^\perp$	a local cycle in $\mathcal{S}^\perp$	S
$\mathsf{cyc}^\parallel(\mathbf{x})$	a cycle in $\mathcal{S}$ through $\mathbf{x}$	
$\mathsf{cyc}(\sigma)$	the quasi-local cycle at σ	$\hat{\mathbf{r}}, \sigma$
$K_S(\mathbf{x})$	the direction cone at $\mathbf{x} \in V_S$	$\mathbf{x}, S$
K_σ	the cone at the saddle σ	$\hat{\mathbf{r}}, \sigma$
$\mathsf{contrib}$	set of contributing saddles	$\hat{\mathbf{r}}$
p_σ	the leading spline at σ	$\mathbf{r}, \sigma$
$\mathcal{I}(\sigma)$	the quasi-local vanishing ideal	σ
Λ	Hessian at a nondegenerate critical point	$\mathbf{r}$, σ

Let $\mathbf{b} = (b_{ij})$ denote the matrix of coefficients of the linear polynomials $\hat{l}_j = 1 - l_j$, and let $\tilde{\mathbf{b}}_j := (b_{1j}x_1, \ldots, b_{dj}x_d)^T$ denote the normal to the real hypersurface $\{l_j(e^{x_1}, \ldots, e^{x_d}) = 0\}$. By $\det_\sigma(\mathbf{b})$ we denote the volume of the zonotope that the rows of $\mathbf{b}$ indexed by $S(\sigma)$ has in $\mathcal{S}_S^\perp$.

Let $\mathcal{A}$ be the intersection lattice of the affine arrangement $\{l_1, \ldots, l_n\}$ in $\mathbb{C}^d$. Formally, this is the lattice of all maximal subsets $S \subseteq \{1, \ldots, n\}$ having the same intersections: $S \in \mathcal{A}$ if and only if no $l_k, k \notin S$ vanishes on the variety defined by $\{l_j, j \in S\}$. We order the lattice by inclusion, not by reverse inclusion as is sometimes done, and we omit the empty set and the whole space. A *sub-arrangement* $\mathcal{A}' \preceq \mathcal{A}$ is the intersection lattice of any subset of these hyperplanes.

Associated with each $S \in \mathcal{A}$ are its dimension, $\dim S$, a flat V_S, a stratum $\mathcal{S}_S$ and a set of saddles $\Sigma(S)$. The flat V_S is simply the variety defined by $\{l_j : j \in S\}$ (thus by definition of S, $V_{S \cup \{k\}} \neq V_S$ for $k \notin S$). The dimension $\dim S$ is the dimension of V_S, and is a grading function for $\mathcal{A}$. We use $\operatorname{codim} S$ for $d - \dim S$. Define

$$\mathcal{S}_S := V_S \setminus \left(\bigcup_{T < S} V_T \right) . \tag{3.1}$$

For each stratum $\mathcal{S}_S$ and each $\mathbf{r} \in (R^+)^d$, define the set of saddles $\Sigma(S)$ of $\mathcal{S}_S$ with respect to the direction $\mathbf{r}$ as the set of $\mathbf{x} \in V_S$ for which the gradient $(-r_1/x_1, \ldots, -r_d/x_d)$ of f at $\mathbf{x}$ is orthogonal to V_S. These are critical values of f_{V_S}. The function f is convex on each orthant, whence in each orthant, either f_{V_S} is not bounded below or f_{V_S} has a unique critical point which is a minimum. We conclude that the cardinality of σ_S intersected with each orthant is 0 or 1 and that the set

$$\Sigma := \bigcup_{S \in \mathcal{A}} \Sigma(S)$$

is the set of minima of f restricted to the connected components of V_S in orthants of $\mathbb{R}^d$. We denote by $S(\sigma)$ the $S \in \mathcal{A}$ for which $\sigma \in \Sigma(S)$. We write $\dim \sigma$ for $\dim S(\sigma)$ and we write $|\sigma|$ for $\sum_{j \in S(\sigma)} n_j$, that is for the degree of the pole of F at σ. Figure 3 shows an arrangement with two one-dimensional strata and one zero-dimensional stratum. One of the one-dimensional strata has no saddle (lack of a saddle in a stratum happens only non-generically, when the corresponding flat is parallel to a coordinate hyperplane). Often every stratum has exactly one saddle, though Examples 1 and 2 show this need not be the case.

The critical points depend on $\mathbf{r}$ only through $\hat{\mathbf{r}}$. For each of these points, σ, it turns out to be important whether $\sigma \in \mathcal{S}_S$, the other possibility being $\sigma(S) \in \mathcal{S}_T$ for some $T < S$. Define the set non of non-generic directions to be the set of $\mathbf{r}$ such that for some σ, $\sigma \in \mathcal{S}_T$ for some $T < S(\sigma)$. The complement, gen, of non is the set of $\mathbf{r}$ such that $\sigma \in \mathcal{S}_{S(\sigma)}$ for all $\sigma \in \Sigma$. The set non is a central hyperplane arrangement, that is, a union of hyperplanes through the origin (cf. Proposition 6.1) and we often consider non and gen to be in $(\mathbb{RP}^{d-1})^+$.

The closure of each stratum $\mathcal{S}$ is an affine subspace, so it has an orthogonal complement $\mathcal{S}_S^\perp$. We summarize the genericity definitions from before as follows, noting that (3.3) is in the hypotheses of all our asymptotic results.

$$\textbf{genericity of } \mathcal{A} \qquad S \in \mathcal{A} \Leftrightarrow |S| \leq d. \tag{3.2}$$

$$\textbf{genericity of } \mathbf{r} \qquad \text{For each } S \in \mathcal{A}, \Sigma(S) \subseteq \mathcal{S}_S. \tag{3.3}$$

The columns $\mathbf{b}_j$ of $\mathbf{b}$ for $j \in S$ span the normal $\mathcal{S}_S^\perp$ to the stratum $\mathcal{S}_S$ (or to the flat V_S). Let $\mathsf{high} := \max_{\sigma\in\Sigma} f_{\hat{\mathbf{r}}}(\sigma)$, and let $\mathsf{low} := \min_{\sigma\in\Sigma} f_{\hat{\mathbf{r}}}(\sigma)$.

Definition 3 (fundamental cycles for generic arrangements) *Under genericity of $\mathcal{A}$ (3.2), define the normal link $\mathsf{cyc}^\perp(\sigma) \subseteq \mathcal{S}_{S(\sigma)}^\perp$, which depends on σ only through $S(\sigma)$, to be the product of circles around zero in each of these coordinates:*

$$\mathsf{cyc}^\perp(S) := \{\mathbf{x} : |\hat{l}_j(x)| = \epsilon \ \forall j \in S\} .$$

Here we think of ϵ as infinitesimal, taking it to be any sufficiently small number, and we use $\hat{l}_j$ to denote the homogenized linear function l_j, i.e., recentered to vanish at the origin. For $\mathbf{x} \in V_S$, define $\mathsf{cyc}^\parallel(\mathbf{x}) := V_S \cap (\mathbf{x} + i\mathbb{R}^d)$ so that if $\sigma \in \Sigma$, then

$$\mathsf{cyc}^\parallel(\sigma) := V_{S(\sigma)} \cap (\sigma + i\mathbb{R}^d)$$

is the imaginary fiber of S passing through the saddle σ. We then define the quasi-local cycle

$$\mathsf{cyc}(\sigma) := \mathsf{cyc}^\perp(S(\sigma)) \times \mathsf{cyc}^\parallel(\sigma),$$

which is the set of $\mathbf{x}+\mathbf{y}$ with $\mathbf{x} \in \mathsf{cyc}^\perp$ and $\mathbf{y} \in \mathsf{cyc}^\parallel$.

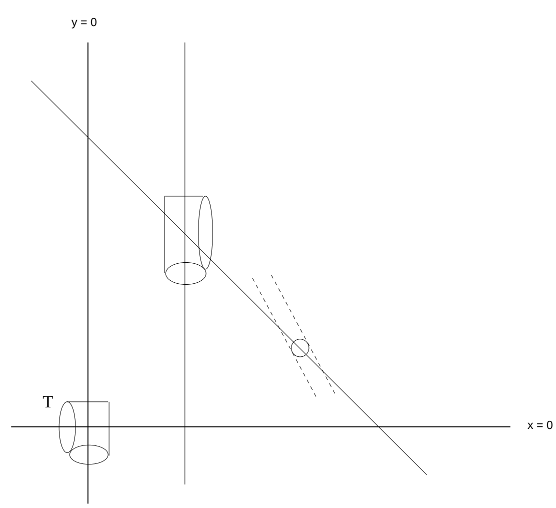

Figure 1: An arrangement with two saddles

Example 1 *Figure 3 shows a two-dimensional arrangement with two strata of co-dimension 1, whose flats intersect in a stratum of co-dimension 2.* **Warning: only the two real dimensions of the four-dimensional space $\mathbb{C}^2$ are drawn.** *The torus T is shown (taking liberties with the third and fourth dimensions), as well as the local 2-cycle $\mathsf{cyc} = \mathsf{cyc}^\perp$ at the zero-dimensional stratum. One of the one-dimensional strata has no saddle (the vertical one), while the other does have a saddle, σ. The local cycle $\mathsf{cyc}^\perp(\sigma)$ is a circle, and $\mathsf{cyc}^\|(\sigma)$ is a line (in an imaginary direction); thus $\mathsf{cyc}(\sigma)$ is an infinite cylinder.*

Example 2 *Consider the arrangement of planes $3x + y + 4z = 8$ and $2x + y + 2z = 5$ meeting in the line $x = 3 - 2z, y = 2z - 1$. The two-dimensional strata each have a unique saddle, but the one-dimensional stratum intersects the positive orthant and the orthant $x, z > 0, y < 0$ both in bounded segments, each containing a saddle. This example and the last show that a stratum may have no saddles or may have two or more saddles.*

To complete the definition we need to choose an orientation for $\mathsf{cyc}(\sigma)$ (we will not need to choose orientations for $\mathsf{cyc}^\perp$ and $\mathsf{cyc}^\|$ separately). A convenient choice is as follows. Let S denote $S(\sigma)$. Removing the real hyperplanes $\{V_j \cap \mathbb{R}^d : j \in S\}$ from $\mathbb{R}^d$ leaves $2^{|S|}$ open orthants, precisely one of which, call it B, contains the origin. Each orthant R is a d-chain in the complement of $\bigcup_{j\in S} V_j$ in either $\mathbb{R}^d$ or $\mathbb{C}^d$, and as the latter, has a well defined signed intersection number with the d-cycle $\mathsf{cyc}(\sigma)$ (cf. further discussion of the dual basis of signed intersection numbers in Section 4.2). The cardinality of the intersection is always 1, so the signed intersection numbers take on values of ± 1 as R is reflected in each V_j. We choose an orientation of $\mathsf{cyc}(\sigma)$ so that the signed intersection number is $+1$ on B. For later use, we remark that this choice is natural with respect to inclusion of the complement of an arrangement in the complement of a sub-arrangement.

The following example, which is a specific case of the general example analyzed in [BM93, Example 4.5], serves to illustrate the above definitions.

Example 3 (two complex lines in two space) *Let $d = n = 2$, let $3l_1 = 3 - x - 2y$, let $3l_2 = 3 - 2x - y$, and let $F = 1/(l_1 l_2)$, that is, $\eta = 1$.*

The two lines intersect at $p = (1, 1)$. The strata are:

$$\begin{aligned} \mathcal{S}_{12} &= \{p\} ; \\ \mathcal{S}_1 &= V_1 \setminus \{p\} = \{l_1 = 0 \neq l_2\} ; \\ \mathcal{S}_1 &= V_2 \setminus \{p\} = \{l_2 = 0 \neq l_1\} . \end{aligned}$$

If $\hat{\mathbf{r}} = (\alpha, 1 - \alpha)$ then the saddles are:

$$\begin{aligned} \sigma_{12} &= \{p\} ; \\ \sigma_1 &= \left(3\alpha, 3\frac{1-\alpha}{2}\right) ; \\ \sigma_2 &= \left(3\frac{1-\alpha}{2}, 3\alpha\right) . \end{aligned}$$

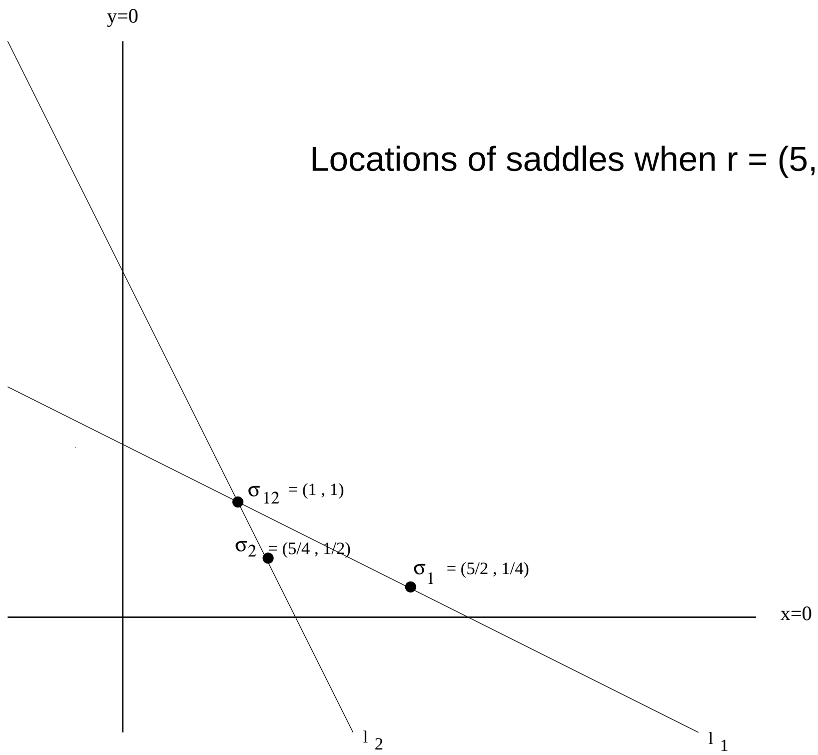

Figure 2: a worked example

The normal to $\mathcal{S}_1$ is the complex line $\mathcal{S}_1^\perp = \{(z,2z)\}$ while $\mathcal{S}_2^\perp = \{(2z,z)\}$, and $\mathcal{S}_{12}^\perp$ is all of $\mathbb{C}^2$. The cycle $\mathsf{cyc}^\parallel(1)$ is the set

$$\{(3\alpha + 2it, 3\frac{1-\alpha}{2} - it) : t \in \mathbb{R}\}$$

while $\mathsf{cyc}^\parallel(12)$ is the single point p. The cycle $\mathsf{cyc}^\perp(1)$ is the circle $\epsilon e^{i\theta}(1,2)$, where ϵ is any fixed sufficiently small positive number and $\theta \in [0,2\pi)$. Taking the product with $\sigma(1) + \mathsf{cyc}^\parallel(1)$, we see that $\mathsf{cyc}(1)$ is an infinite cylinder. The cycle $\mathsf{cyc}^\perp(12)$ is equal to $\mathsf{cyc}(12)$ and is a small torus $|3-x-2y| = |3-2x-y| = 3\epsilon$ near the point $(1,1)$.

Given $\mathbf{x} \in V_S$, define its direction cone

$$K_S(\mathbf{x}) = (\mathbb{R}^+)^d \cap \mathrm{Pos}\,\{\tilde{\mathbf{b}}_j : j \in S\} \tag{3.4}$$

to be the intersection of the positive orthant of $\mathbb{R}^d$ with the positive hull of the vectors $\tilde{\mathbf{b}}_j$ for $j \in S$. In other words, $K_S(\mathbf{x})$ is the set of positive normals to V_S at $\mathbf{x}$ in logarithmic coordinates.

The cone $K_\sigma := K_{S(\sigma)}(\sigma)$ turns out to be the set of vectors $\mathbf{r}$ for which the integral in Step 3 yields a nonzero contribution, as long as η is not in the ideal $\mathcal{I}(\sigma)$. Thus we are led to define a subset $\mathsf{contrib} \subseteq \Sigma$, depending on $\mathbf{r}$, as follows.

Definition 4 *Let* $\mathsf{contrib}$ *be the set of* $\sigma \in \Sigma$ *for which* $\mathbf{r} \in K_\sigma$.

We remark that $\mathbf{1}_{\mathsf{contrib}}$ is continuous on gen, that is, it is constant on chambers and that it is equivalent to use the relative interior of K_σ in place of K_σ in the above definition because $\mathbf{r} \in \mathsf{non}$ whenever $\mathbf{r}$ is in the relative boundary of K_σ (cf. Proposition 6.1).

Proposition 3.1 *Let* f *be a strictly convex function on* $\mathbb{R}^d$ *and* B *an intersection of halfspaces* $\{\mathbf{y} : \mathbf{y} \cdot \mathbf{v}_j \leq 0\}_{1 \leq j \leq k}$ *in* $\mathbb{R}^d$. *Let* $\mathbf{x} \in \partial B$ *and order the halfspaces so that for some* $1 \leq k_0 \leq k$, $\mathbf{x} \cdot \mathbf{v}_j = 0$ *exactly when* $j \leq k_0$. *Then* f *has a minimum on* B *at* $\mathbf{x}$ *if and only if* $-\nabla f \in \operatorname{Pos}\{\mathbf{v}_j : j \leq k_0\}$.

PROOF: Assume without loss of generatlity that $\mathbf{x}$ is the origin. If $-\nabla f$ is in the positive hull, then for every $\mathbf{x}$ with $\epsilon\mathbf{x} \in B$ for some $\epsilon > 0$, we have $-\nabla f \cdot \mathbf{x} \leq 0$. By strict convexity of f, this implies that $f(\lambda\mathbf{x}) > f(0)$ for every $\lambda > 0$, which implies that f has a unique minimum on B at 0. Conversely, if $-\nabla f$ is not in the positive hull of $\{\mathbf{v}_j : j \leq k_0\}$ then $-\nabla f \cdot \mathbf{x} > 0$ for some $\mathbf{x} \in B$, whence for some $\epsilon > 0$, $f(\epsilon\mathbf{x}) < f(0)$, and f is not minimized on B at 0. $\square$

Applying this to the alternate interpretation of K_S, with $f = f_{\mathbf{r}}$, we see that $\sigma \in \mathsf{contrib}$ if and only if σ is the location of the minimum of f on an intersection of halfspaces bounded by the hyperplanes containing $S(\sigma)$:

$$\sigma \in \mathsf{contrib} \Longleftrightarrow f(\sigma) = \min\{f(\mathbf{x}) : \sum_{i=1}^d b_{ij} x_i \leq 1 \, \forall j \in S(\sigma)\} \, . \tag{3.5}$$

Example 3 continued: *Continuing Example 3, we see that in case* $K_p(p)$ *is the cone between slopes* $1/2$ *and 2. The positive arc of* $\mathbb{RP}^1$ *is broken into three segments, between the slopes 0 and* $1/2$, *between* $1/2$ *and 2, and between 2 and* $+\infty$. *Over each of these,* σ_1 *and* σ_2 *are both not equal to p, that is, both genericity assumptions hold. The set* $\mathsf{contrib}$ *always contains all saddles in strata of codimension 1. Since* $K_p(p)$ *is the middle arc, we have that* $\sigma_{12} \in \mathsf{contrib}$ *if and only if* $\mathbf{r}$ *is in the middle arc (see Figure 3).*

Example 4 *Consider the generating function whose pole set is drawn in Figure 4:*

$$F(\mathbf{z}) := \frac{2}{(1-2x)(1-2y)} \cdot \left[2 + \frac{xy - 1}{1 - xy(1 + x + y + 2xy)}\right]$$

The diagonal of this function generates the number of small balls necessary to cover the equilibrium measure of a Markov chain used in [Kar02] to model a TCP-IP protocol. The denominator of F *contains three factors,* $(1-2x)$, $(1-2y)$ *and* $(1 - xy(1+x+y+2xy))$. *The last of these is nonlinear, but the same analysis applies. For* $\hat{\mathbf{r}} = (1,1)$ *(the diagonal direction), all saddles are in* $\mathsf{contrib}$. *However, on the one where* f *is maximized, namely the intersection point* $(1/2, 1/2)$ *of the two lines, the function* η *vanishes; this implies*

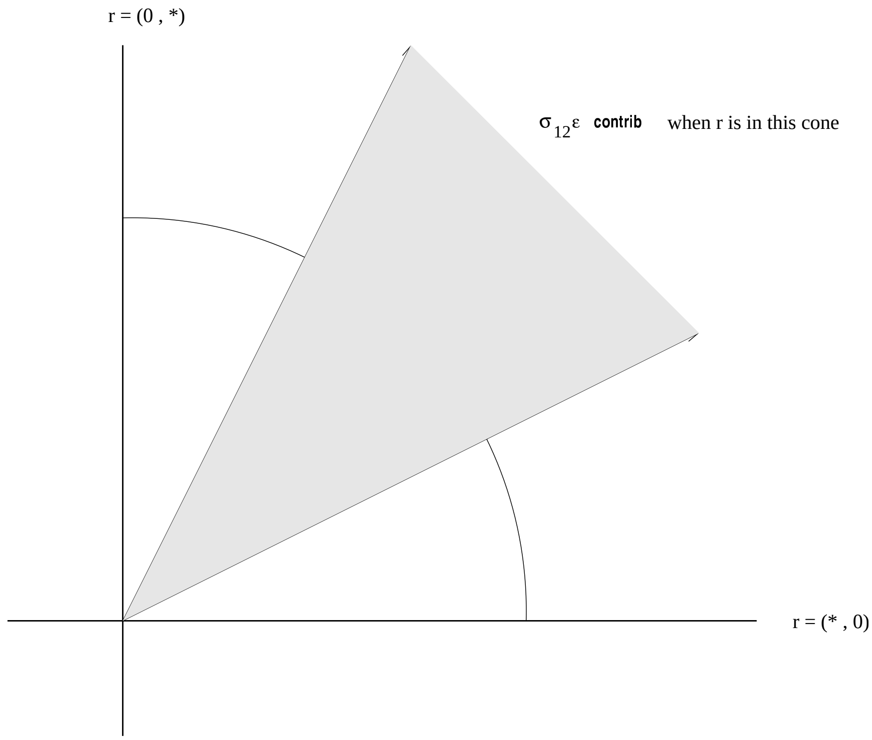

Figure 3: Cones in **r**-space

$\eta \in \mathcal{I}(1/2, 1/2)$ so that the contribution here vanishes (cf. step 4e of Algorithm 1.1). The leading term comes from the next highest saddles, which are the other two intersection points (there is a symmetry in x and y), which dominate the saddles of dimension one and yield the leading asymptotics. This example is included to demonstrate that in relevant examples, one might find that η vanishes on the highest saddle in contrib *and therefore need to look at lower saddles, which is not possible with the technology in [PW04].*

Having explained the loop where the leading saddle is identified, we turn finally to the computation of the integral there. Given $S \in \mathcal{A}$ and a point $\mathbf{x} \in \mathcal{S}_S$, the function $F \prod_{j \in S} l_j$ is holomorphic in a neighborhood of $\mathbf{x}$. Let $\mathcal{S}_S^\perp$ denote the affine space through $\mathbf{x}$ orthogonal to $\mathcal{S}_S$.

Definition 5 (the polynomial p_σ) *Given a point $\mathbf{x}$ in a stratum $\mathcal{S}_S$, let Y be the positive orthant of a real vector space with n_j coordinates for each $j \in S$. Define a linear map $\phi_{S,\mathbf{x}} : Y \to K_S(\mathbf{x})$ by sending each of the n_j standard basis vectors associated with $j \in S$ to the vector $\mathbf{b}_j$. Let λ^k denote Lebesgue measure on $(\mathbb{R}^+)^k$ and define*

$$p_\sigma := (2\pi i)^{\operatorname{codim} \sigma} \cdot \text{density of } \phi_{S(\sigma),\sigma}[\lambda^{|\sigma|}] \tag{3.6}$$

to be the density of the image under ϕ of Lebesgue measure on $(\mathbb{R}^+)^{|\sigma|}$. It is evident that p_σ is piecewise polynomial of degree $|\sigma| - \operatorname{codim} \sigma$; it is in fact a spline.

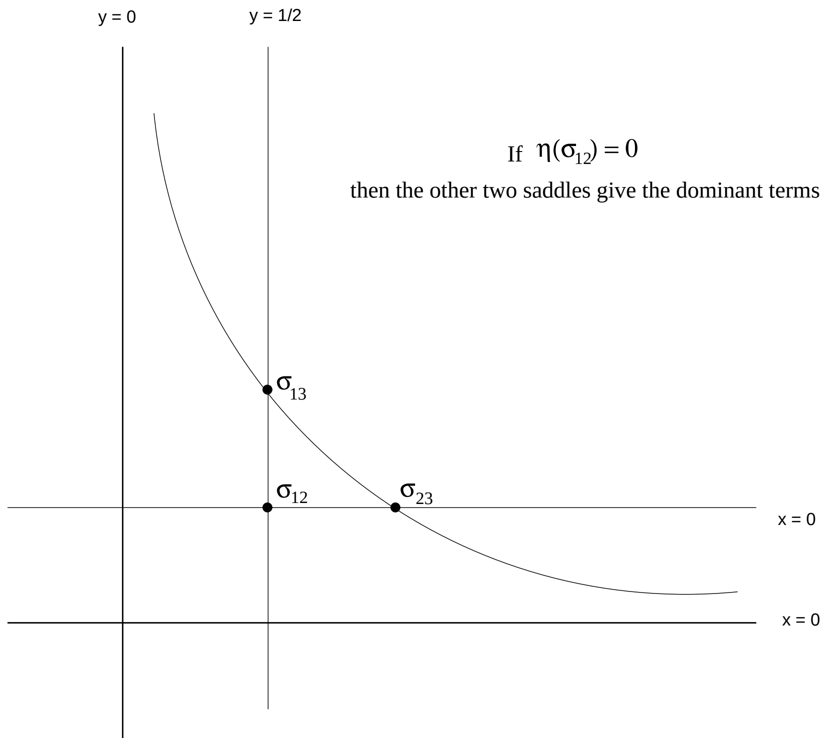

Figure 4: similar methods as for a hyperplane arrangement

Definition 6 (the polynomial P_σ) *Denote by Φ_σ the differential operator $p_\sigma[\partial]$ one gets by formally substituting $(-\partial/\partial x_i)$ for x_i in p_σ. Let $\eta_\sigma := \eta / \prod_{j \notin S(\sigma)} l_j^{n_j}$. Define*

$$P_\sigma := P_{\sigma,\mathbf{z}} := \Phi_\sigma(\mathbf{z}^{-\mathbf{r}-\mathbf{1}} \eta_\sigma) . \tag{3.7}$$

Under genericity of $\mathbf{r}$ (3.3), η_σ is analytic in a neighborhood of σ, so P_σ varies analytically with $\mathbf{z}$. One sees as well from its form that P_σ is a polynomial in $\mathbf{r}$. When $\eta_\sigma(\mathbf{z}) \neq 0$, the leading term of this polynomial is gotten by taking all partial derivatives in $\mathbf{z}^{-\mathbf{r}-\mathbf{1}}$ rather than in η_σ. The k^{th} negative partial derivative $\partial/\partial x_j$ multiplies by a factor of $(r_j + k)/x_j$. If we are only interested in the leading term in $\mathbf{r}$, we may approximate this by r_j/x_j, meaning that

$$p_\sigma[-\partial](\mathbf{z}^{-\mathbf{r}-\mathbf{1}} \eta_\sigma)(\sigma) \sim p_\sigma(r_1/x_1, \ldots, r_d/x_d) \cdot \sigma^{-\mathbf{r}} \eta_\sigma(\sigma) . \tag{3.8}$$

In the case $\sigma = \mathbf{1}$, this simplifies to

$$P_{\sigma,\mathbf{z}} = p_\sigma(\mathbf{r}) \left(\eta_\sigma(\mathbf{1}) + O(|\mathbf{r}|^{-1}) \right) .$$

Example 3 continued further: *the Jacobian of ϕ_p is the constant function $1/3$, so that $p_p = P_p \equiv 3$. In general, p_σ has degree zero whenever $V_{S(\sigma)}$ is a transverse intersection; to be independent of $\mathbf{z}$ as well is* *unusual. We compute the normal integral for $\sigma_{12} = p$:*

$$\det{}_{12}(\mathbf{b}) = \det(\mathbf{b}) = b_{11}b_{22} - b_{12}b_{21}$$

we then see from Theorem 2.2 that (since $p^{-\mathbf{r}-\mathbf{1}} \equiv 1$):

$$\int_{\mathsf{cyc}(12)} \xi = 3(2\pi i)^2 .$$

For any $\mathbf{r}$, the function $f_{\mathbf{r}}$ is nonpositive on Σ. Thus when $\mathbf{r}$ is in the middle arc, since $\eta(p) \neq 0$, we have $\sigma_{12} \in \mathsf{contrib}$ and the dominant contribution to the sum in Theorem 2.1 is from σ_{12}:

$$a_{\mathbf{r}} = 3 + \sum_{\sigma \neq p} \int_{\mathsf{cyc}(\sigma)} \xi = \frac{1}{3} + O\left(\exp(|\mathbf{r}| \max_{i=1,2} f_{\hat{\mathbf{r}}}(\sigma_i))\right) .$$

When $\mathbf{r}$ is in the upper arc, or if η is replaced by a function vanishing at p, then the leading contribution to the sum in Theorem 2.1 comes from σ_1 (respectively σ_2 for the lower arc). From Theorem 2.2 we see that

$$a_{\mathbf{r}} = \int_{\mathsf{cyc}(1)} \frac{1}{6\pi i} \int_{\mathsf{cyc}(1)} \mathbf{z}^{-\mathbf{r}-\mathbf{1}} \frac{\eta(\mathbf{z})}{l_2(\mathbf{z})} + O\left(\exp(f_{\mathbf{r}}(\sigma_2))\right) \tag{3.9}$$

which has a saddle point expansion as $\sigma(1)^{-\mathbf{r}-\mathbf{1}}$ times a series of terms of order $\mathbf{r}^{-j-1/2}$ for $j = 0, 1, 2, \ldots$.

When $\mathcal{A}$ is generic, it is easy to determine the highest contributing saddle: after computing $\mathsf{contrib}$, one eliminates saddles σ where η vanishes (cf. the last statement in Theorem 2.2). In general, the elimination of saddles requires that an ideal membership question be settled, the effectiveness of which depends on the nature of the decscription of η. One must choose a ring to work in. Let $\mathcal{R}$ be any ring containing $\mathbb{C}[z_1, \ldots, z_d]$ and containing η. Instead of eliminating saddles where η vanishes, we eliminate saddles σ for which $\Phi_\sigma(\mathbf{z}^{-\mathbf{r}-\mathbf{1}}\eta_\sigma)$ is identically zero in a neighborhood of σ in $\mathsf{cyc}^{\|}(\sigma)$, the ideal of such functions eing denoted below by $\mathcal{I}(\sigma)$.

Two final definitions occurring in Theorem 2.5 are the saddle point determinant at a quadratically nondegenerate saddle and the vanishing ideal.

Definition 7 *Let $\sigma = \sigma(S)$ be a critical point for the function $f_{\mathbf{r}}$ restricted to V_S. Suppose that* $\dim(S) > 0$ *and parametrize the real subspace of V_S by the orthonormal vectors $y_1, \ldots, y_k$ as in the definition of $\mathbf{z}^{\|}$. Then the Hessian (matrix of second partial derivatives) of $f_{\mathbf{r}}$ with respect to $y_1, \ldots, y_k$ has a determinant which we denote $\Lambda_S(\mathbf{r})$, which is a homogeneous function of $\mathbf{r}$ of degree k.*

Definition 8 (vanishing ideal) *Given the critical point σ in the stratum S, define the ideal*

$$\mathcal{I}(\sigma) = \left\langle \prod_{j \in S'} l_j^{n_j} : \dim\langle \mathbf{b}_j : j \notin S' \rangle < \operatorname{codim} S \right\rangle$$

to be the ideal in the ring of germs of analytic functions at σ of functions η' for which the pole set of $\eta' / \prod l_i^{n_i}$ does not have S as a flat. Observe that if S is a single point and $n_j = 1$ for $j \in S$, then this is the ideal of all functions vanishing at σ; if S is a singleton $\{j\}$ then $I(\sigma)$ is generated by $l_j^{n_j}$.

In order to make the last few definitions less abstract, we give three more examples.

Example 3 completed: *We continue from (3.9) in the case that* $\mathbf{r}$ *is in the upper arc, so* $\sigma^* = \sigma_1$. *Since* S_1 *has co-dimension 1, the polynomial* p_{σ^*} *is just the constant*

$$|\mathbf{b}|^{-1} = \left|(\frac{1}{3}, \frac{2}{3})\right|^{-1} = \frac{3}{\sqrt{5}} .$$

Recall that $\sigma_1 = (3\alpha, 3(1-\alpha)/2)$ *where* $r_1/r_2 = \alpha/(1-\alpha) < 1/2$. *We then have* $\eta/l_2 = 2/(3-9\alpha)$. *To evaluate* Λ_{σ^*}, *according to Theorem 2.5, we need to compute the second derivative of* $-r_1 \log x_1 - r_2 \log x_2$ *with respect to arclength as* $\mathbf{x}$ *varies along a line that may be conveniently parametrized as* $\sigma_1 + t(1,2)/\sqrt{5}$. *This comes out to be*

$$\Lambda_1 = \frac{4}{45}\left(\frac{r_1}{\alpha^2} + \frac{r_2}{(1-\alpha)^2}\right) .$$

Putting this all together gives

$$\begin{aligned} a_{\mathbf{r}} &= (1+o(1))\frac{3}{\sqrt{5}}\frac{2}{3-9\alpha}\left(2\pi\frac{4}{45}(\frac{\alpha|\mathbf{r}|}{\alpha^2} + \frac{(1-\alpha)|\mathbf{r}|}{(1-\alpha)^2}\right)^{-1/2}(3\alpha)^{-r_1}(\frac{3}{2}(1-\alpha))^{-r_2} \\ &= \frac{3}{1-3\alpha}\sqrt{\frac{\alpha(1-\alpha)}{2\pi}}|\mathbf{r}|^{-1/2}\exp\left\{|\mathbf{r}|\left[\alpha\log(3\alpha) + (1-\alpha)\log((3/2)(1-\alpha))\right]\right\} . \end{aligned}$$

Observe that the asymptotics cannot be uniform near $\alpha = 1/3$ *since the denominator* $1-3\alpha$ *of the asymptotic expression vanishes.*

Example 5 (three concurrent lines) *To Example 3, add a third linear factor to the denominator,* $3l_3 = 3 - (3/2)x - (3/2)y$. *This also passes through* p, *so the lattice* $\mathcal{A}$ *is non-generic, having* p *at the bottom, covered by each of three one-dimensional strata. The cone* K_p *is the same as before, but the set of* $\tilde{\mathbf{b}}$ *vectors, contains not only* $(1,2)$ *and* $(2,1)$ *but also* $(3/2, 3/2)$. *Thus* p_p *is now a spline of degree 1, the image of Lebesgue density under a map sending the three standard basis vectors to* $(1,2), (2,1)$ *and* $(3/2, 3/2)$ *respectively. This comes out to the tent function* $2\min\{r,s\} - \max\{r,s\}$.

The local ring of germs of functions at p *has a unique maximal ideal* $\mathfrak{m}$, *namely those functions vanishing at* p. *In the previous cases* p_σ *had degree zero, so vanishing of* $\Phi_\sigma(\mathbf{z}^{-\mathbf{r}-\mathbf{1}}\eta_\sigma)(\mathbf{z}^{\|})$ *was equivalent to the vanishing of* η *on* $\mathbf{z}^{\|}$. *In the present case,* $\mathcal{I}(\sigma)$ *is generated by* $l_1 l_2$, $l_1 l_3$ *and* $l_2 l_3$; *these are linearly dependent so it suffices to take* $l_1 l_2$ *and* $l_2 l_3$, *which in turn generate the ideal* $\mathfrak{m}^2$ *of germs vanishing to order two at* σ_1. *The quotient by this ideal then has dimension three, with coset representatives* $1, l_1$ *and* l_3; *thus* $\eta/\prod l_j$ *may be written as*

$$\frac{A}{l_1 l_2 l_3} + \frac{B}{l_1 l_2} + \frac{C}{l_2 l_3} + \frac{\eta_*}{l_1 l_2 l_3}$$

with $\eta_* \in \mathcal{I}(\sigma)$. *These four terms generate coefficient arrays that are respectively a multiple of the tent function, a constant in one sub-chamber of* K_p, *a constant on the other sub-chamber of* K_p, *and an array with exponentially decaying magnitudes.*

Example 6 (not all $b_{ij} \geq 0$**)** *In the previous examples,* $b_{ij} \geq 0$, *which implies that all coefficients* $a_{\mathbf{r}}$ *are nonnegative. Automatically, in that case, the highest contributing saddle is minimal (in the lattice order) among all contributing saddles. Thus in those cases, the saddle at* p *was always dominant when it* *contributed. To see what can happen when this assumption is removed, consider an example where $d = n = 2$, with $2l_1 = 2 - x - y$ and $l_2 = 1 - (10/9)y + (1/9)x$. When $\hat{\mathbf{r}} = (\alpha, 1-\alpha)$, the saddles are given by*

$$\begin{aligned} \sigma_1 &= (2\alpha, 2(1-\alpha)) \\ \sigma_2 &= (-9\alpha, \frac{9}{10}\alpha) \\ p := \sigma_{12} &= (1,1)\,. \end{aligned}$$

Ordering these by height we find that

$$\begin{array}{lll} \sigma_2 > p > \sigma_1 & \textit{on} & 0 < \alpha < 0.317\ldots \\ p > \sigma_2 > \sigma_1 & \textit{on} & 0.317\ldots < \alpha < \dfrac{\log(20/9)}{\log 10} \approx 0.346\ldots \\ p > \sigma_1 > \sigma_2 & \textit{on} & \dfrac{\log(20/9)}{\log 10} < \alpha < \dfrac{1}{2} \\ p = \sigma_1 > \sigma_2 & \textit{at} & \alpha = 1/2 \\ p > \sigma_1 > \sigma_2 & \textit{on} & \alpha > 1/2\,. \end{array}$$

The saddles σ_1 and σ_2 are always in $\mathsf{contrib}$, but $p \in \mathsf{contrib}$ only when $\alpha < 1/2$, or equivalently, when $\mathbf{r}$ is above the diagonal. Thus for small α we have exponentially large $a_{\mathbf{r}}$, but there is a constant added in which takes over when the exponentially large expression turns into an exponentially decaying expression as α increases past a point approximately equal to $0.317\ldots$. As α nears $1/2$, another exponentially decaying term becomes nearly constant, reaching a maximum of $\Theta(r^{-1/2})$ on the diagonal. This term takes over abruptly as the constant term ceases to contribute below the diagonal. This transition is like the one that occurs in Example 3.

4 Description of homology and proofs of Theorems 2.1, 2.2 and 2.3

In this section we collect some results on the topology of $(\mathcal{M}, -\infty)$ and use them to prove the first three theorems, that is, the ones not relying on cohomological computations. We will draw on some standard results from Stratified Morse theory, which we quote from [GM88]. We prove a few more lemmas that are specific to the present analysis, leaving to an appendix some of the more technical geometric steps. Homology, here and throughout, refers to homology of the singular chain complex.

4.1 Relative and local topology

Subsequent analysis will show that the integral (1.2) is composed of summands, indexed by critical points σ of $f_{\mathbf{r}}$, each of which is $O(\exp(|\mathbf{r}| f_{\hat{\mathbf{r}}}(\sigma)))$ where σ is a critical point of f. Any quantity that is $O(e^{f(x)})$ for some $x \in \mathcal{M}_{\mathsf{low}}$ is negligible in comparison. With this in mind, the following proposition explains why we work in relative homology. The proof is given in the appendix.

Proposition 4.1 *Let ω be any d-form of at most polynomial growth and holomorphic outside of an affine set,* poles*. Let* low *be the least critical value of $f_{\mathbf{r}}$ on* poles *and let $b <$* low*.*

1. *The integral $\int_C \mathbf{z}^{-\mathbf{r}}\omega$ is well defined as a function of $[C] \in H_d(\mathcal{M}, \mathcal{M}_b)$ up to terms that are $O(e^{-|\mathbf{r}|b})$ as $\mathbf{r} \to \infty$.*
2. *The integral $\int_C \mathbf{z}^{-\mathbf{r}}\omega$ is well defined as a function of $[C] \in H_d(\mathcal{M}, -\infty)$.*

Definition 9 (imaginary fibers) *For $\mathbf{x} \in \mathbb{R}^d$, let $\gamma_{\mathbf{x}}$ denote the d-cycle $\{\mathbf{x} + i\mathbf{y} : \mathbf{y} \in \mathbb{R}^d\}$. This respects strata, meaning that if $\mathbf{x} \in \mathcal{M}$ then $\gamma_{\mathbf{x}} \subseteq \mathcal{M}$, and if $\mathbf{x} \notin V_S$ then $\gamma \cap V_S = \emptyset$.*

Definition 10 (normal links) *For $S \in \mathcal{A}$ assume $\sigma_S \in \mathcal{S}_S$ and recall that $\mathcal{S}_S^{\perp}$ is the orthogonal complement to the linear (not affine) space $\mathcal{S}_S$. Let $\mathcal{M}_S^{\perp}$ denote the link of the hyperplane arrangement in the normal slice, that is, the complement of the union of those flats V_U containing V_S in $\mathcal{S}_S^{\perp}$.*

Proposition 4.2 *If the central arrangement for the normal slice is generic (that is, if $|S| = \operatorname{codim} S$) then $H_{\operatorname{codim} S}(\mathcal{M}_S^{\perp})$ is cyclic and generated by the cycle $\mathsf{cyc}^{\perp}(S)$ defined in Definition 3. If the arrangement is not generic then $H_{\operatorname{codim} S}(\mathcal{M}_S^{\perp})$ is a free abelian group and we denote its rank by $n(S)$.*

Proof: The first statement is immediate from the fact that $\mathcal{M}_S^{\perp}$ is homeomorphic to a product of punctured disks. The second follows from the fact that complements of complex affine varieties in d dimensions are homotopy equivalent to CW-complexes of dimension at most d. $\square$

The main lemma we rely on from Stratified Morse Theory, also proved in the appendix, is the decomposition of $H_d(\mathcal{M}, -\infty)$ into a direct sum of local homology groups at each saddle; these groups will be denoted $H_{d,\sigma}(\mathcal{M})$. There are a number of equivalent definitions of these: homology relative to a small neighborhood of the complement, homology of $\mathcal{M}_b$ relative to $\mathcal{M}_a$ where (a, b) contains $f(\sigma)$, homology of the complement of the sub-arrangement corresponding to σ, or the homology of the poset of $\mathcal{A}$ above σ localized to the top element. For clarity, we take $H_{d,\sigma}$ to be an abstract space and define the following maps.

Observe that $\mathcal{M}_S^{\perp}$ retracts to an arbitrarily small neighborhood of the origin, so any class in $H_{\operatorname{codim} S}(\mathcal{M}_S^{\perp})$ has a *local* cycle representative supported on an arbitrarily small neighborhood of the origin. For such a cycle, C, we recall from Definition 3 that for $\mathbf{x} \in V_S$,

$$i_{\mathbf{x}}([C]) := \left[C \times \mathsf{cyc}^{\parallel}(\mathbf{x})\right]$$

denotes the class of the product set $\{\mathbf{v}+\mathbf{y} : \mathbf{v} \in C, \mathbf{y} \in \mathsf{cyc}^{\parallel}(\mathbf{x})\}$, with $\mathsf{cyc}^{\parallel}(\mathbf{x}) = (\mathbf{x} + i\,\mathbb{R}^d) \cap V_S$.

Lemma 4.3 (quasi-local decomposition) *Assume genericity of* $\mathbf{r}$. *For each* $\sigma \in \Sigma$ *there is a group* $H_{d,\sigma}(\mathcal{M})$ *and a natural surjection* $\pi_\sigma : H_d(\mathcal{M}, -\infty) \to H_{d,\sigma}(\mathcal{M})$ *having the following properties.*

1. $\Phi_\sigma := \pi_\sigma \circ i_\sigma$ *is a natural isomorphism between* $H_{\mathrm{codim}\,\sigma}(\mathcal{M}^{\perp}_{S(\sigma)})$ *and* $H_{d,\sigma}(\mathcal{M})$.
2. $\bigoplus_{\sigma\in\Sigma} \pi_\sigma$ *is an isomorphism between* $H_d(\mathcal{M}, -\infty)$ *and* $\bigoplus_{\sigma\in\Sigma} H_{d,\sigma}(\mathcal{M})$.

Naturality of the maps is with respect to sub-arrangements in the following sense. If $\mathcal{A}'$ *is a sub-arrangement of* $\mathcal{A}$ *then the inclusion of* $\mathcal{M}$ *in* $\mathcal{M}'$ *(complementation reverses the direction) induces a map* ι_* *on homology which commutes with the above maps. Specifically, if* S *is a flat of both* $\mathcal{M}$ *and* $\mathcal{M}'$ *and* $\gamma \in H_d(\mathcal{M}, -\infty)$ *then for* $\sigma \in \Sigma(S)$,

$$\iota_*(\pi_\sigma(\alpha)) = \pi_\sigma(\iota_*(\alpha))$$

while if S *is a flat of* $\mathcal{M}$ *but not a flat of* $\mathcal{M}'$ *then for* $\sigma \in \Sigma(S)$,

$$\iota_*(\pi_\sigma(\alpha)) = 0\,.$$

$$\begin{tikzcd}
H_{\mathrm{codim}\,\sigma}(\mathcal{M}^{\perp}_{S(\sigma)}) \arrow[r, "i_\sigma"] \arrow[dr, "\Phi_\sigma"', "\cong"] & H_d(\mathcal{M}, -\infty) \arrow[d, "\pi_\sigma"] \\
& H_{d,\sigma}(\mathcal{M})
\end{tikzcd}$$

Remark: The fact that $\bigoplus_\sigma \pi_\sigma$ is an isomorphism is always true of top-dimensional homology of any stratified space, since it comes from the construction of $\mathcal{M}_c$ up to homotopy equivalence by the attachment of d-dimensional CW complexes as c passes critical values. The fact that there is a map i_σ allowing us to go back is special to the case of hyperplane arrangements, where relative cycles in $(\mathcal{M}_b, \mathcal{M}_a)$ may always be extended to cycles in $(\mathcal{M}, -\infty)$.

4.2 Another basis for $H_d(\mathcal{M}, -\infty)$

The next step in proving Theorem 2.1 is to identify the class of T with respect to another basis.

Definition 11 (fiber basis) *Let $\mathcal{B}$ denote the set of bounded components B of $\mathbb{R}^d \cap \mathcal{M}$. Let $\mathsf{bd} = \bigcup_{B \in \mathcal{B}} \overline{B}$ denote the union of the closures of the bounded regions. Clearly, if $\mathbf{x}$ and $\mathbf{x}'$ are in the same component of $\mathbb{R}^d \cap \mathcal{M}$ then $\gamma_{\mathbf{x}} = \gamma_{\mathbf{x}'}$ in $H_d(\mathcal{M}, -\infty)$, and if $\mathbf{x}$ is in an unbounded component of $\mathbb{R}^d \cap \mathcal{M}$ then $\gamma_{\mathbf{x}} = 0$ in $H_d(\mathcal{M}, -\infty)$. For $B \in \mathcal{B}$, let γ_B denote the d-cycle $\gamma_{\mathbf{x}}$ for some $\mathbf{x} \in B$.*

Proposition 4.4 *The cycles γ_B are a basis for $H_d(\mathcal{M}, -\infty)$ as B varies over $\mathcal{B}$.*

One proof, via an explicit retraction of the complement of bd to $-\infty$, is given in the appendix. Since we will need several descriptions of the homology of $(\mathcal{M}, -\infty)$, we give here a different proof of Proposition 4.4. The following definition relies on our assumption of genericity of $\mathbf{r}$.

Definition 12 *Define a map $\alpha : \mathcal{B} \to \Sigma$ by $\alpha(B) = \sigma \in \Sigma$ such that $f|_{\overline{B}}$ is minimized at σ.*

To see that the minimum really occurs at a saddle, note that $\overline{B}$ is a polytope whose faces are subsets of the flats V_S. The minimum on a face must occur at a critical point for f or on the boundary. Thus the minimum of f, which is unique by convexity, occurs in the relative interior of some face, hence is the critical point for the corresponding stratum.

Lemma 4.5 *Assuming genericity of $\mathbf{r}$, $|\mathcal{B}| \geq |\mathcal{A}|$, and in fact α is surjective with*

$$|\alpha^{-1}(\sigma)| = \dim H_{d,\sigma}(\mathcal{M}) .$$

If the arrangement is generic (3.2), then the map α is one to one, whence $|\mathcal{B}| = |\Sigma|$.

Proof: Given σ and $\mathbf{r}$, the normals $\mathbf{b}_j$ to the hyperplanes $\{V_j : j \in S(\sigma)\}$ span a space containing $-\nabla f(\sigma)$. Thus $-\nabla f(\sigma)$ is in the positive hull of $\{\epsilon_j \mathbf{b}_j : j \in S(\sigma)\}$ for some choice of signs $\epsilon \in \{\pm 1\}^S$. The intersection of halfspaces $\{\mathbf{x} : \mathbf{x} \cdot \epsilon_j \mathbf{b}_j \leq 0\}$ contains a unique component B of $\mathbb{R}^d \setminus (\mathsf{poles} \cup \mathsf{planes})$ with $\sigma \in \overline{B}$, on which, by Proposition 3.1, f is minimized. Since $f \to -\infty$ at infinity, we see that B is bounded, i.e., $B \in \mathcal{B}$ and $\alpha(B) = \sigma$. Thus α is surjective.

If the arrangement is generic then the normals $\{\mathbf{b}_j : j \in S\}$ are independent, so there is a unique choice of ϵ for which $-\nabla f \in \mathrm{Pos}\,\{\epsilon_j \mathbf{b}_j : j \in S\}$, and it follows again from Proposition 3.1 that $\alpha^{-1}(\sigma)$ is unique. For the proof that dimension is counted by α^{-1} in the case of a non-generic arrangement, see tha appendix. $\square$

Proof of Proposition 4.4: For $B \in \mathcal{B}$, let I_B denote the map on $H_d(\mathcal{M}, -\infty)$ taking any homology class to its signed intersection number with the set B. Since $\partial B \subseteq \mathsf{poles} \cup \mathsf{planes}$, it is immediate to verify that this is well defined on homology classes. It is also evident that $I_B(\gamma_{B'}) = \delta_{B,B'}$. By Lemma 4.3 and Lemma 4.5, the cardinality of $\mathcal{B}$ is equal to the dimension of $H_d(\mathcal{M}, -\infty)$. It follows that $\{\gamma_B : B \in \mathcal{B}\}$ is a basis for $H_d(\mathcal{M}, -\infty)$ and that $\{I_B : B \in \mathcal{B}\}$ is a dual basis. $\square$

4.3 Proofs of Theorems 2.1, 2.2 and 2.3

The basis $\{\gamma_B : B \in \mathcal{B}\}$ is good for computation since for any cycle C we have

$$C = \sum_{B \in \mathcal{B}} I_B(C)\gamma_B \, .$$

A good beginning for expressing $[T]$ in the $\mathsf{cyc}(\sigma)$ basis is to write each in the γ_B basis. Since T is the product of small oriented circles in each coordinate, we see that

$$I_B(T) = \begin{cases} 0 & \text{if } 0 \notin \overline{B}, \\ \mathbf{sgn}(B) & \text{if } 0 \in \overline{B}, \end{cases} \tag{4.1}$$

where $\mathbf{sgn}(B) = \mathbf{sgn}(\mathbf{x})$ for any $\mathbf{x} \in B$. On the other hand, by choice of orientation, decomposing $\mathsf{cyc}(\sigma)$ in the $\{\gamma_B\}$ basis requires

$$I_B(\mathsf{cyc}(\sigma)) = \begin{cases} 0 & \text{if } \sigma(\sigma) \notin \overline{B}, \\ \mathbf{sgn}(B, \sigma) & \text{if } \sigma \in \overline{B}, \end{cases} \tag{4.2}$$

where $\mathbf{sgn}(B, \sigma)$ is $(-1)^N$ if N hyperplanes containing σ separate B from the origin.

Figure 4.3 depicts, for the arrangement in Example 3, how the identity we are trying to prove, namely $T = \sum_{\sigma \in \mathsf{contrib}} \mathsf{cyc}(\sigma)$, looks in the γ_B basis. The term $\mathbf{sgn}(\sigma)$ is not present since all saddles are in the positive orthant. The $+$ and $-$ terms in the regions are the values of $\mathbf{sgn}(B, \sigma)$. These must add up to 0 in any bounded region not touching the origin, and must add up to 1 for the region that does touch the origin. The figure shows how these identities are preserved as $\mathbf{r}$ varies: as $\mathbf{r}$ crosses the non-generic value where $\sigma_2 = \sigma_{12}$, the value of $\mathbf{1}_{\mathsf{contrib}}(\sigma_{12})$ changes from 0 to 1. Both of the other saddles have co-dimension 1 so are always in $\mathsf{contrib}$. We now prove that something like this works in general.

Proof of Theorem 2.1: We rely on the following reformulation of Proposition 4.4:

> Let $H^c_d(\mathcal{M} - \infty)$ be the homology with closed support of of $(\mathcal{M}, -\infty)$. Then the cycles corresponding to the (standardly oriented) bounded components $B \in \mathcal{B}$ form a basis in $H^c_d(\mathcal{M}, -\infty)$, and the pairing $H_d(\mathcal{M}, -\infty) \otimes H^c_d(\mathcal{M}, -\infty) \to \mathbb{Z}$ given by the algebraic intersection number is non-degenerate.

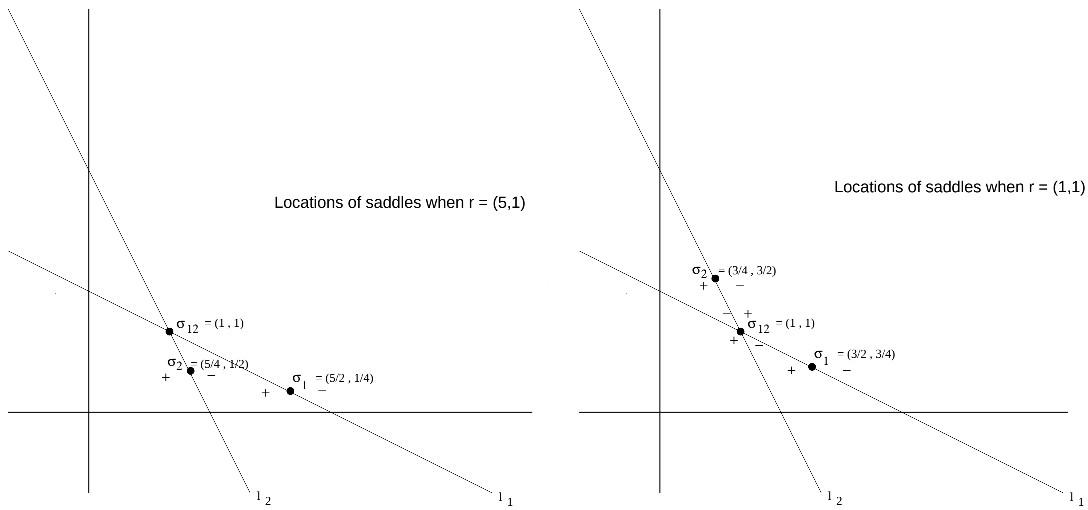

In particular, to check Theorem 2.1 it is enough to verify that the intersection numbers of right- and left-hand sides in 2.3 with any of the cycles spanned by bounded components in $\mathcal{B}$ is the same. The intersection number of $[T]$ with $[B]$ (i.e. the cycle corresponding to $B \in \mathcal{B}$) is easy to compute: it is $\mathbf{sgn}(\prod_i x_i), x \in B$ if closure of B contains the origin, and 0 otherwise. To prove that this is the same as the intersection number of the RHS of 2.3 with B, we employ a deformation trick. Specifically, we construct a 1-parametric family of functions $\{\phi_s : s \in \mathbb{R}^+\}$ such that

- Generically, the functions $\phi_{\cdot}$ are Morse;
- For large s, the sum of indices of critical points of ϕ_s in B is the same as the intersection index of $[T]$ with $[B]$;
- For small s, the sum of indices of critical points of ϕ_s in B is the same as the intersection index of $\sum_{\sigma \in \mathsf{contrib}} \mathbf{sgn}(\sigma)\, \mathsf{cyc}(\sigma)$ with $[B]$, and
- The critical points of ϕ_s remain in a compact subset of B when s varies in a compact subset of $\mathbb{R}^+$.

The last property ensures the conservation of the total index of ϕ_s, which by the two previous properties, then implies the desired equality.

We construct the family with the desired properties for the positive orthant $(\mathbb{R}^+)^d$ only; the proof extends to other orthants immediately. Let

$$\phi_s = f_{\mathbf{r}} + s \sum_j \frac{1}{l_j} .$$

For *all* s, the restriction of ψ_s to the component B_o in $\mathcal{B}$ adjoining the origin is strictly convex, and therefore has a unique critical point there. As for the other bounded components, let us denote by

$$R := \sup_{\mathbf{x}\in(\mathbb{R}^+)^d} \frac{\inf_{T\mathbf{x}\in(\mathbb{R}^+)^d\setminus\mathsf{bd}} T}{\sup_{t\mathbf{x}\in B_o} t}$$

the largest ratio of proportional vectors lying in B_o and in an unbounded component, correspondingly. Then, the restriction of ψ_s to the line through a vector is, up to an additive constant,

$$\psi_s(t\mathbf{x}) = -r\log t + s\sum_j \frac{a_j}{a_j - t}, \tag{4.3}$$

where $r = \sum r_i$ and $0 < a_1 < \ldots < a_n \leq Ra_1$. Now, to show that for s large enough, there is no critical point of ϕ_s in the bounded components in $\mathcal{B}$ other than B_o, it is enough to show that (4.3) has no critical points between the poles a_1 and a_n. This is easy: taking the derivatives yields

$$-\frac{r}{t} + s\sum_j \frac{a_j}{(a_j - t)^2} \geq -\frac{r}{t} + s\frac{a_1}{(a_1 - t)^2},$$

and an elementary computation shows that for $s > rR^2$ the right hand side of the inequality above is positive in the interval (a_1, a_n).

We turn now to small s. Outside of small vicinity of the poles (shrinking with s), ψ_s is C^1-close to $f_{\mathbf{r}}$, and therefore has no critical points.

Let σ be a critical point of restriction of $f_{\mathbf{r}}$ to a k-codimensional stratum V. As we assumed genericity of the arrangement, V is the zero locus of k functions $l_i, i \in I, |I| = k$. Fix a coordinate system $\xi = (\xi_k)_{k=1,\ldots,d}$ centered at σ in which $\{l_i\}, i \in I$ are given by $\xi_i, i = 1, \ldots, k$. Then

$$\phi_s = \sum_{i=1}^{k} a_i\xi_i + q(\xi) + s\left(\sum_{i=1}^{k}\frac{1}{\xi_i} + c(\xi)\right), \tag{4.4}$$

where a_i are partial derivatives of $f_{\mathbf{r}}$ with respect to ξ_i's. Clearly, q and c are holomorphic in a vicinity $U \ni \sigma$ (which we assume to not intersect any of the other poles) and $dq|_\sigma = 0$.

Assume for instance genericity of $\mathbf{r}$. This implies that all a_i are nonzero. A point $\xi \in U$ is critical for ϕ_s if

$$a_i\xi_i^2 + (q_i + sc_i)\xi_i^2 = s \quad \text{for} \quad i = 1, \ldots, k \tag{4.5}$$

and

$$q_i = -sc_i \quad \text{for} \quad i = k+1, \ldots, d \tag{4.6}$$

(here c_i and q_i stand for partial derivatives of c and q with respect to ξ_i).

One sees immediately, that if at least one of coefficients a_i is negative, there are no real solutions to (4.5) for s small enough. On the other hand, if all $a_i, i = 1, \ldots, k$ are positive, there exist 2^k

solutions of (4.5,4.6) depending ϕ_s which tend, as $s \to 0$, to σ: to see this one can consider 2-covering $s = t^2$ and look for solutions $\xi_i = t\xi'(t)$. A straightforward application of the implicit function theorem implies existence of 2^k solutions ξ' for t small enough.

The constructed critical points ϕ_s belong to different components into which hyperplanes $\{x_i = 0\}, i = 1, \ldots, k$ split U. The second derivative of ϕ_s at a critical point ξ is

$$H_{\phi_s} = s^{-1/2}\text{diag}(h_1(s), h_2(s), \ldots, h_k(s), 0, \ldots, 0) + H_q + sH_c, \tag{4.7}$$

where H_{ϕ_s}, H_q and H_c are the Hesse matrices of corresponding functions, and h_i are Puiseux series in $s^{1/2}$ with constant terms $h_i(0) = \mathbf{sgn}(\xi_i)a_i^{3/2}$.

For small s the index of the critical point ξ in a component of U is $(-1)^l$, where l is the number of negative coordinates ξ_i on the component. On the other hand, this is exactly the index of the intersection of $\mathsf{cyc}(\sigma)$ with the component containing ξ, which implies the result. $\square$

Proof of Theorem 2.2: Begin by writing $\int_{\mathsf{cyc}(\sigma)} \xi$ as

$$\int_{\mathsf{cyc}^\parallel} d\mathbf{z}^\parallel \int_{\mathsf{cyc}^\perp} \xi \, d\mathbf{z}^\perp . \tag{4.8}$$

Letting $y_j = l_j(\mathbf{z})$, for $j \in S$, we have that $d\mathbf{y} = \det{}_\sigma(\mathbf{b}) \, d\mathbf{z}^\perp$. Coordinatizing each $\mathbf{z} \in \mathbb{C}^d$ by $\mathbf{z}^\parallel$ and $\mathbf{y}$, we may integrate

$$\int_{\mathsf{cyc}^\perp} \xi \, d\mathbf{z}^\perp = \det{}_\sigma(\mathbf{b})^{-1} \int_U \mathbf{z}^{-\mathbf{r}-\mathbf{1}} \frac{\eta(\mathbf{z})}{\prod_{j \notin S} l_j(\mathbf{z})} \frac{d\mathbf{y}}{\mathbf{y}}$$

where U is the torus $\{|y_j| = \epsilon, j \in S\}$. Iteratively, one may integrate each dy_j/y_j to get $2\pi i$ times the value at $y_j = 0$, and one arrives at

$$(2\pi i)^{\text{codim}\,\sigma} \det{}_\sigma(\mathbf{b})^{-1} \mathbf{z}^{-\mathbf{r}-\mathbf{1}} \frac{\eta(\mathbf{z})}{\prod_{j \notin S(\sigma)} l_j(\mathbf{z})}$$

for the inner integral in (4.8). This proves the theorem. $\square$

Proof of Theorem 2.3: Fix a saddle σ. The quasi-local decomposition lemma gives us $\mathsf{cyc}(\sigma)$. What is left to show is that this is nonzero if and only if $\sigma \in \mathsf{contrib}$.

By definition, $\sigma \in \mathsf{contrib}$ if and only if $\mathbf{r}$ is in the positive hull of $\tilde{\mathbf{b}}_j$ for $j \in S(\sigma)$. Since we have assumed $\mathbf{r} \notin \partial K_S(\sigma)$, we see that this holds if and only if

$$\mathbf{r} \text{ is in the positive hull of } \{\tilde{\mathbf{b}}_j : j \in U\} \text{ for some } U \subseteq S(\sigma) \text{ of cardinality codim}\,(\sigma)\,. \tag{4.9}$$

For such a U, let $\mathcal{A}_U$ denote the generic sub-arrangement of $\mathcal{A}$ whose elements are the subsets of U and let $\mathcal{M}_U$ denote the complement. If we now consider the collection of inclusions $\mathcal{M} \subseteq \mathcal{M}_U$, we see that $\sigma \in \mathsf{contrib}$ if and only if $\sigma \in \mathsf{contrib}$ in the arrangement $\mathcal{A}_U$ for some U. By the naturality clause in the quasi-local decomposition lemma, $\mathsf{cyc}(\sigma)$ maps under these inclusions to the quasi-local cycles $\mathsf{cyc}_U(\sigma)$ in the decomposition of $H_d(\mathcal{M}_U, -\infty)$. By Theorem 2.1 and (4.9), if $\sigma \in \mathsf{contrib}$ then at least one of these images is nonzero, hence $[T] \neq 0$ in $H_{d,\sigma}(\mathcal{M}, -\infty)$. On the other hand, if $\sigma \notin \mathsf{contrib}$ then all the images are zero. We shall see in the next section that the images of the cohomology generators for $H^{d,\sigma}(\mathcal{M}_U)$ generate $H^{d,\sigma}(\mathcal{M})$, so that the vanishing of all the images of $\mathsf{cyc}(\sigma)$ in $H_{d,\sigma}(\mathcal{M}, -\infty)$ implies the vanishing of $\mathsf{cyc}(\sigma)$. □

Together with equation (1.2) and part (ii) of Proposition 4.1, this yields the following corollary:

Corollary 4.6 *For sufficiently large* $\mathbf{r}$*, there is an exact equality*

$$a_{\mathbf{r}} = \frac{1}{(2\pi i)^d} \sum_{\sigma \in \mathsf{contrib}} \int_{\mathsf{cyc}(\sigma)} \mathbf{sgn}(\sigma) \xi \, .$$

□

5 Remaining proofs

5.1 Bases for cohomology

We want to build on our understanding of $H_{d,\sigma}(\mathcal{M}) \cong H_{\mathrm{codim}\,\sigma}(\mathcal{M}^{\perp}_{S(\sigma)})$ enough to integrate ξ over classes in this group. We need therefore to examine the dual group, which we call $H^{d,\sigma}(\mathcal{M})$, in order to determine which element of the dual corresponds to integration against ξ. In order to understand the dual to $H_{\mathrm{codim}\,(S)}(\mathcal{M}^{\perp}_{S(\sigma)})$, we examine central arrangements with a single point of common intersection.

Suppose j hyperplanes $\{l_i = 0\}_{1 \leq i \leq j}$ in $\mathbb{C}^d$, $d \leq j$, have a single point, say the origin, as their common intersection. The d-dimensional cohomology $H^d(X)$ of the complement X of this arrangement is generated by forms

$$\omega_Q := \prod_{i \in Q} \frac{1}{l_i} \, d\mathbf{z}$$

as Q ranges over sets of cardinality d such that the $\{l_i : i \in Q\}$ are independent. One may choose as a basis the forms ω_Q where Q is in the *no broken circuit complex*, BC, defined to be those Q containing no broken circuit; here, a circuit of the matroid of independence of the linear forms $\{l_i : 1 \leq i \leq j\}$ is a minimal dependent set, and a broken circuit is a circuit with its greatest element deleted. If $j = d$, then the arrangement is generic, there is only one set of cardinality d, so the d-dimensional homology is cyclic and generated by $\omega_{\{1,\ldots,d\}}$. These facts follow from [OT92, Theorems 3.43, 3.126 and 5.89]. Specifically, it is proved that BC indexes the Orlik-Solomon algebra, which is isomorphic to another algebra, which is isomorphic to $H^*(X)$.

The d-dimensional homology of X is generated by local cycles. Integration against local cycles is a well defined operation on germs of d-forms near the origin. Thus we have a simple proposition extending the utility of the $\{\omega_Q\}$ basis.

Proposition 5.1 *Suppose $j = d$. Let G be a function analytic in a neighborhood of the origin in X. Then*

$$[G \cdot \omega_Q] = [G(0) \cdot \omega_Q] \quad \textit{in } H^d(X)\,.$$

PROOF: Suppose $G(0) = 0$. The functions vanishing at 0 are generated over germs of analytic functions by $\{l_i : i \in Q\}$. Thus $G = \sum_{i \in Q} l_i G_i$ and so

$$G \cdot \omega_Q = \sum_{i \in Q} \frac{G_i}{\prod_{j \neq i} l_j}\,.$$

The i^{th} summand ω_i in this sum is a generator for the cohomology of the complement of the sub-arrangement of the $d-1$ hyperplanes other than $\{l_i = 0\}$. Letting ι be the inclusion of X into the complement of this arrangement, we see that $\iota^*(\omega_i) = \iota^*(0) = 0$ since the complement of the sub-arrangement has no d-dimensional cohomology. Thus $G \cdot \omega_Q = \sum_i \omega_i = 0$ in $H^d(X)$, and the proposition follows from linearity. $\square$

Returning to the non-generic case $j \geq d$, we resolve local holomorphic d-forms on X into the $\{\omega_Q\}$ basis in several steps.

Let $\mathbf{n}$ be a vector of length j of nonnegative integers and write $\mathbf{l}^{-\mathbf{n}}$ for $\prod_{i=1}^{j} l_i^{-n_i}$. Say the support of the monomial $\mathbf{l}^{-\mathbf{n}}$ is the set of i for which $n_i \neq 0$.

Algorithm 5.2

1. *Initialize W to be $\mathbf{l}^{-\mathbf{n}}$.*
2. *Repeat until W is empty:*
   (a) *If a monomial in W of the form $c\,l^{-\mathbf{m}}$ contains no broken circuit among its support then output it and remove it from W.*
   (b) *If a monomial of the form $c\,l^{-\mathbf{m}}$ contains a broken circuit U among its support then let k be greater than every element of U with $U \cup \{k\}$ a circuit and choose c_i not all zero so that $l_k = \sum_{i \in U} c_i l_i = 0$. Replace $c\,l^{-\mathbf{m}}$ in W by*

$$\sum_{i \in U} c\,c_i\,l^{-\mathbf{m}+\delta_i-\delta_k}\,.$$

At each replacement step, a monomial is replaced by a sum equal to the original monomial. Also, at each replacement step a monomial is replaced by a sum of monomials with smaller weight, where the weight of a monomial $l^{-\mathbf{m}}$ is $\sum i\, m_i$. Since weights well order monomials, the algorithm must halt. The output is therefore a way of writing $\mathbf{l}^{-\mathbf{n}}$ as the sum of monomials containing no broken circuit in their support. If the support of $\mathbf{l}^{-\mathbf{n}}$ is large enough to contain at least one element of BC then each monomial will have this property. We may generalize Algorithm 5.2 to forms $G \cdot \mathbf{l}^{-\mathbf{n}}\, d\mathbf{z}$ by operating only on the $\mathbf{l}^{-\mathbf{n}}$ part.

Next, pick nonzero vectors $\{\mathbf{v}(Q,i) : Q \in \mathsf{BC}, i \in Q\}$ so that each $l_k(\mathbf{v}(Q,i)) = 0$ for all $k \in Q$ with $k \neq i$. These vectors are uniquely defined up to scalar multiples. Let $G \cdot \mathbf{l}^{-\mathbf{n}}\, d\mathbf{z}$ be a d-form with support in $Q \in \mathsf{BC}$. We use the symbol $(\frac{\partial}{\partial l_i})_Q$ to denote the differential operator on locally analytic functions defined by

$$\left(\frac{\partial}{\partial l_i}\right)_Q (G) := \frac{\nabla G \cdot \mathbf{v}(Q,i)}{-\mathbf{b}_i \cdot (\mathbf{v}(Q,i))}\,.$$

Lemma 5.3 *For any locally analytic G, any monomial $\mathbf{l}^{-\mathbf{n}}$ with support in $Q \in \mathsf{BC}$ and any i in the support of $\mathbf{l}^{-\mathbf{n}}$,*

$$\left[G \cdot \mathbf{l}^{-\mathbf{n}-\delta_i}\, d\mathbf{z}\right] = \frac{1}{n_i}\left(\frac{\partial}{\partial l_i}\right)_Q (G) \cdot \mathbf{l}^{-\mathbf{n}}\, d\mathbf{z}$$

in $H^d(X)$.

PROOF: For any orthonormal basis $\{\mathbf{x}^1, \ldots, \mathbf{x}^d\}$ the differential of any form $\Psi := \psi \cdot \bigwedge_{k\geq 2} dx_k$ is equal to

$$(\partial\psi/\partial\mathbf{x}^1) \cdot dV$$

where $\{dx_k\}$ is dual to $\{(d/d\mathbf{x}^k)\}$ and $dV := \bigwedge_{k\geq 1} dx_k$. Choosing a basis with $\mathbf{x}^1 = \mathbf{v}(Q,i)$ and observing that $\partial l_k/\partial\mathbf{x}^1 = 0$ for $k \in Q \setminus \{i\}$ gives

$$d\Psi = (\partial\psi/\partial\mathbf{v}(Q,i))\, dV.$$

The differential $d\Psi$, being exact, is zero in $H^d(X)$. Setting $\psi = G \cdot \mathbf{l}^{-\mathbf{n}}$ now gives

$$\left[\frac{\partial G}{\partial\mathbf{v}(Q,i)}\mathbf{l}^{-\mathbf{n}}\, dV - n_i\frac{\partial l_i}{\partial\mathbf{v}(Q,i)} G\mathbf{l}^{-\mathbf{n}-\delta_i}\right] dV = 0$$

in $H^d(X)$, whence

$$\left[\frac{\partial G}{\partial\mathbf{v}(Q,i)}\mathbf{l}^{-\mathbf{n}}\right] = n_i\left[\frac{\partial l_i}{\partial\mathbf{v}(Q,i)} G\,\mathbf{l}^{-\mathbf{n}-\delta_i}\right]$$

in $H^d(X)$. To get the result, divide by $n_i\partial l_i/\partial\mathbf{v}(Q,i)$ and observe that $\partial l_i/\partial\mathbf{v}(Q,i) = -\mathbf{b}_i \cdot \mathbf{v}(Q,i)$. $\square$

Corollary 5.4 (cohomology representation) *If $\mathbf{l}^{-\mathbf{n}}$ has support precisely Q then*

$$\left[G \cdot \mathbf{l}^{-\mathbf{n}}\, d\mathbf{z}\right] = \frac{1}{\prod_{i\in Q}(n_i - 1)!}\prod_{i\in Q}\left(\frac{\partial}{\partial l_i}\right)_Q^{n_i-1} (G)(\mathbf{0}) \cdot \omega_Q\,.$$

PROOF: Apply the lemma inductively to lower $\mathbf{n}$ to a 0-1 vector with the same support, then apply Proposition 5.1. $\square$

5.2 Proof of Theorem 2.4

Now we return to the setting of an arbitrary hyperplane arrangement $\mathcal{A}$ and the form ξ which is the integrand of (1.2). Fix a $\sigma \in \Sigma$ and a $\mathbf{z}^{\|} \in \mathsf{cyc}^{\|}(\sigma)$. On $\mathsf{cyc}^{\perp}(\mathbf{z}^{\|})$, which is defined to be $\mathbf{z}^{\|} + \mathsf{cyc}^{\perp}(\sigma)$, the form ξ restricts to a form which we also, without ambiguity, can denote ξ. We use the notation $\mathsf{BC}(\sigma)$ to denote the set of $U \subseteq S(\sigma)$ containing no broken circuit. The space $X = \mathcal{M}^{\perp}_{S(\sigma)}$ is identified with $\mathsf{cyc}^{\perp}(\mathbf{z}^{\|})$, with the origin moving to $\mathbf{z}^{\|}$.

We start by writing

$$\int_T \xi = \sum_{\sigma \in \Sigma} \int_{\mathsf{cyc}(\sigma)} \xi = \sum_{\sigma \in \Sigma} \int_{\mathsf{cyc}^{\|}(\sigma)} \left[\int_{\mathsf{cyc}^{\perp}} \xi \, d\mathbf{z}^{\perp} \right] d\mathbf{z}^{\|} .$$

In each inner integral, we may now replace ξ by anything equal to ξ in the dual space $H^{d,\sigma}$ to the local homology group $H_{d,\sigma}$. In other words, we may compute in separate dual spaces for each σ. Since η_σ is holomorphic in a neighborhood of the stratum $S(\sigma)$, we may apply Algorithm 5.2, or specifically the generalization to forms $G \cdot \mathbf{l}^{-\mathbf{n}} \, d\mathbf{z}$, separately at each σ. We then use Corollary 5.4 to produce a cohomology representative in $H^{d,\sigma}$, namely

$$[\xi] = \left[\sum_{Q \in \mathsf{BC}(\sigma)} \alpha(Q, \mathbf{r}, \mathbf{z}) \, (\omega_Q \wedge d\mathbf{z}^{\|}) \right] . \tag{5.10}$$

We remark that after applying Algorithm 5.2 the coefficients are not constant on fibers of $\mathbf{z} \mapsto \mathbf{z}^{\|}$, but after applying Corollary 5.4, the coefficients $\alpha(Q, \mathbf{r}, \mathbf{z})$ depend on $\mathbf{z}$ only through $\mathbf{z}^{\|}$.

Replacing ξ by the LHS of (5.10) in each inner integral, we sum to get

$$\int_T \xi = \sum_{\sigma \in \Sigma} \int_{\mathsf{cyc}(\sigma)} I_{\xi,\sigma}(\mathbf{z}^{\|}) \, d\mathbf{z}^{\|} , \tag{5.11}$$

where

$$I_{\xi,\sigma}(\mathbf{z}^{\|}) := \sum_{Q \in \mathsf{BC}(\sigma)} \alpha(Q, \mathbf{r}, \mathbf{z}^{\|}) \, \mathbf{sgn}(Q) \, \mathbf{1}_{\sigma \in \mathsf{contrib}_Q} \, |\det Q|^{-1} ,$$

and is zero whenever $\sigma \notin \mathsf{contrib}$.

PROOF OF THEOREM 2.4: Part 1 of Theorem 2.4 has only to do with properties of the functions p_σ and P_σ, and these were noted at the time the definitions were given.

Comparing (5.11) with part 2 of Theorem 2.4, it remains to show that the inner integral $I_{\xi,\sigma}$ is correctly computed by $p_\sigma[-\partial](\mathbf{z}^{-\mathbf{r}-\mathbf{1}}\eta_\sigma)$, or in other words, that

$$I_{\xi,\sigma} = P_\sigma\,. \tag{5.12}$$

case 1: the distinct normals $\{\mathbf{b}_j : j \in S(\sigma)\}$ to the hyperplanes $\{V_j : j \in S(\sigma)\}$ are linearly independent, though they may appear with multiplicities greater than 1. In this case, $\mathsf{BC}(\sigma)$ is the singleton $\{S(\sigma)\}$ and equation (5.11) has only one summand, so that

$$[\xi] = \left[\frac{1}{\prod_{j\in S}(n_j-1)!}\left(\frac{\partial}{\partial \mathbf{l}}\right)_S^{\mathbf{n}-\mathbf{1}}(\eta_\sigma \mathbf{z}^{-\mathbf{r}-\mathbf{1}})\right](\mathbf{z}^{\parallel})\cdot\omega_S\wedge d\mathbf{z}^{\parallel}\,.$$

Thus

$$I(\sigma,\mathbf{z}^{\parallel}) = \mathbf{sgn}(S)\,\mathbf{1}_{\sigma\in\mathsf{contrib}_S}\,|\det S|^{-1}\frac{1}{\prod_{j\in S}(n_j-1)!}\left[\left(\frac{\partial}{\partial \mathbf{l}}\right)_S^{\mathbf{n}-\mathbf{1}}(\eta_\sigma \mathbf{z}^{-\mathbf{r}-\mathbf{1}})\right](\mathbf{z}^{\parallel})\,. \tag{5.13}$$

To complete case 1, we work on an explicit form of P_σ. Renumber the normal vectors $\mathbf{b}_1,\ldots,\mathbf{b}_k$ and denote their multiplicities by $n_1,\ldots,n_k$. The map ϕ in the definition of p_σ is then a map from $\mathbb{R}^{|\sigma|}$ to $\mathbb{R}^k$. Write $\phi=\phi_2\circ\phi_1$ where ϕ_1 maps the n_j standard basis vectors corresponding to the n_j appearances of $\mathbf{b}_j$ all to the j^{th} standard basis vector e_j of $\mathbb{R}_k$, and ϕ_2 maps e_j to $\mathbf{b}_j$. The map of n standard basis vectors all to the standard basis vector in $\mathbb{R}^1$ maps Lebesgue measure on $(\mathbb{R}^+)^n$ to $x^{n-1}(n-1)!\,dx$, as follows from the fact that the preimage of x is the simplex $\{\sum_{j=1}^n x_j = x\}$. Consequently, the density of the image of Lebesgue measure under ϕ_1 is equal to

$$\prod_{j=1}^k \frac{x_j^{n_j-1}}{(n_j-1)!}\,.$$

The map ϕ_2 has constant Jacobian $|\det\mathbf{b}|$, so the density $p_\sigma(\mathbf{y})$ is equal to

$$|\det\mathbf{b}|^{-1}\prod_{j=1}^k\frac{x_j(\mathbf{y})^{n_j-1}}{(n_j-1)!} \tag{5.14}$$

where x_j is the j^{th} coordinate function of the inverse of ϕ_2.

If $\mathbf{v}^j$ is a vector orthogonal to each $\mathbf{b}_i$ for $i\neq j$, then dotting the equation $\mathbf{y}=\sum_{j=1}^k\lambda_j\mathbf{b}_j$ results in $\mathbf{y}\cdot\mathbf{v}^j=\lambda_j\mathbf{b}_j\cdot\mathbf{v}^j$. Thus the function x_j is computed by $(\mathbf{y}\cdot\mathbf{v}^j)/(\mathbf{b}_j\cdot\mathbf{v}^j)$. We plug this into (5.14) to arrive at

$$p_\sigma(\mathbf{y}) = |\det\mathbf{b}|^{-1}\prod_{j=1}^k\left[\frac{1}{(n_j-1)!}\left(\frac{\mathbf{y}\cdot\mathbf{v}^j}{\mathbf{b}_j\cdot\mathbf{v}^j}\right)^{n_j-1}\right].$$

We may substitute $y_j=\partial/\partial x_j$ in this to see that Φ_σ is the operator defined by

$$\Phi_\sigma = |\det\mathbf{b}|^{-1}\prod_{j=1}^k\left[\frac{1}{(n_j-1)!}\left(\frac{\nabla(\cdot)\cdot\mathbf{v}^j}{\mathbf{b}_j\cdot\mathbf{v}^j}\right)^{n_j-1}\right].$$

Going back to Definition 6, we have a function η_σ locally analytic at σ, defined by $\eta = \mathbf{l}^{-\mathbf{n}}\eta_\sigma$ and we have

$$\begin{aligned} P_\sigma &= \Phi_\sigma(\mathbf{z}^{-\mathbf{r}-\mathbf{1}}\eta_\sigma) \\ &= |\det \mathbf{b}|^{-1} \prod_{j=1}^{k} \left[\frac{1}{(n_j-1)!} \left(\frac{\partial}{\partial l_j} \right)^{n_j-1} \right] (\eta_\sigma \mathbf{z}^{\mathbf{r}-\mathbf{1}}) \,. \end{aligned}$$

Comparing this to equation (5.13) proves (5.12) in the special case.

case 2: general ξ. With $\xi = \mathbf{z}^{-\mathbf{r}-\mathbf{1}}\eta_\sigma \mathbf{l}^{-\mathbf{n}}$ and $\mathbf{n}$ having support $S(\sigma)$, we let $s = |\mathbf{n}| = \sum_{i=1}^k n_i$. If $s = \operatorname{codim}\sigma$, then $|S| = \operatorname{codim}\sigma$ and $n_j = 1$ for $j \in S$. We are then in case 1, and we have already proved that $P_\xi = I_{\xi,\sigma}$.

Suppose now for induction that $s > \operatorname{codim}\sigma$ and that the theorem has already been established for $|\mathbf{n}| = s - 1$. In the definition of P_σ, there is an implicit dependence on the support S of the denominator $\mathbf{l}^{\mathbf{n}}$ of ξ; when we compute with ξ^j in place of ξ, this support will be some $Q \in \mathsf{BC}(\sigma)$ and it will be clearer if we write P_{ξ^j} instead of P_σ, where

$$P_{\xi^j} = p_{\xi^j}[-\partial](\mathbf{z}^{-\mathbf{r}-\mathbf{1}}\eta_\sigma)$$

and p_{ξ^j} is the image of the density of $\lambda^{|\mathbf{n}|}$ mapped according to the multiplicity of factors in the denominator of ξ^j. A conclusion of the induction hypothesis is that the correspondence $\xi \mapsto P_\xi$ must satisfy a linearity condition. Specifically, since $P_\xi = I_{\xi,\sigma}$ and $P_{\xi^j} = I_{\xi^j,\sigma}$, and since $I_{\xi^j,\sigma}$ is linear (being an integral), we may conclude that

$$P_\xi = \sum_j P_{\xi^j} \,.$$

Since this is true as a function of $\mathbf{r}$, we conclude that in fact

$$p_\xi = \sum_j p_{\xi_j} \,. \tag{5.15}$$

To complete the induction, we now let ξ have order s and $\xi = \sum_j \xi^j$ as above. Let $m := \max S$ be the largest index in S. We remark that the basis exchange steps in Algorithm 5.2 never remove m from the support of the denominator, so all the terms ξ^j will have $n_m \geq 1$. The form $l_m\xi$ has one fewer factor of l_m in its denominator, so by definition of p_ξ, we have

$$p_\xi(xx) = \int_0^\infty p_{l_m\xi}(\mathbf{x} - \lambda \mathbf{b}_m)\, d\lambda \,. \tag{5.16}$$

The positive cone on $\{\mathbf{b}_j : j \in S\}$ contains no line, and since this cone is closed under addition of positive multiples of $\mathbf{b}_m$, we conclude that $\mathbf{x} - \lambda\mathbf{b}_1$ is not in the cone for any $\mathbf{x}$ in the cone with $\lambda(\mathbf{x})$ sufficiently large. This verifies that upper limit of the integral in (5.16) is finite for each $\mathbf{x}$ and therefore that the integral exists. Equation (5.16) holds as well for any ξ^j in place of ξ. Applying the induction hypothesis to $l_m\xi = \sum_j l_m\xi^j$ we see that

$$\begin{aligned} p_\xi &= \int_0^\infty p_{l_m\xi}(\cdot - \lambda \mathbf{b}_m)\, d\lambda \\ &= \sum_j \int_0^\infty p_{l_m\xi^j}(\cdot - \lambda \mathbf{b}_m)\, d\lambda \\ &= \sum_j p_{\xi^j} \, . \end{aligned}$$

It follows from case 1 and linearity of $I_{\xi,\sigma}$ that

$$\begin{aligned} P_\xi &= \sum_j P_{\xi^j} \\ &= \sum_j I_{\xi^j,\sigma} \\ &= I_{\xi,\sigma} \, , \end{aligned}$$

completing the induction. $\square$

5.3 Saddle integrals

Lemma 5.5 *The function $\phi(\mathbf{z}) := -\hat{\mathbf{r}} \cdot \log \mathbf{z}$ on $\mathcal{S}_S$ has a quadratically nondegenerate critical point at $\sigma(S)$ which is a maximum on $\mathsf{cyc}^\parallel(\sigma(S))$. Consequently, if $\sigma \in \mathsf{contrib}$ and $\eta \notin \mathcal{I}(\sigma)$, then*

$$\int_{\mathsf{cyc}^\parallel(\sigma(S))} \mathbf{z}^{-\mathbf{r}-\mathbf{1}} \frac{\eta(\mathbf{z})}{\prod_{j\notin S} l_j(\mathbf{z})}\, d\mathbf{z}^\parallel$$

has an asymptotic expansion as

$$\frac{\eta}{\prod_{j\notin S} l_j^{n_j}}(\sigma^*) p_{\sigma^*}(\mathbf{r}) \Lambda(\sigma^*)^{-1/2} \sigma^{-\mathbf{r}}$$

times a sequence of decreasing powers of $\mathbf{r}$, whose first k terms may be computed from the partial derivatives of $\eta/\prod_{j\notin S} l_j$ at $\sigma(S)$ up to order k and from the matrix $\mathbf{v}$.

Proof: By strict concavity of the logarithm on $\mathbb{R}$, the function ϕ on $\mathbb{R}^d$ has a single critical point on each bounded component of $\mathcal{S}_S \setminus \mathsf{planes}$, which is a minimum. Complexifying, we find the same critical points, all being minima in the real directions of $\mathcal{S}_S$ and maxima in the imaginary directions. Thus they are maxima on $\mathsf{cyc}^\parallel(\sigma(S))$, this being an imaginary fiber. The Hessian of the logarithm is diagonal in the standard basis and nondegenerate, and hence always quadratically nondegenerate on any affine subspace. The lemma now follows from standard saddle point integration; see, e.g., [BH86, (8.2.63)] for the two variable case or [Won89, Theorem IX.5.3] for the general case. $\square$

Proof of Theorem 2.5: Follows from Theorems 2.3 and 2.4 and Lemma 5.5. $\square$

6 Algorithmic aspects and further discussion

6.1 Chamber decomposition of the set of generic $\mathbf{r}$

As mentioned after the definitions of gen and $\mathsf{contrib}$, the orthant $(\mathbb{R}^+)^d$ is divided into open cones or *chambers* by the removal of the set non of non-generic values of $\mathbf{r}$. We begin by describing this decomposition.

Proposition 6.1 *For each flat S of $\mathcal{A}$ and each flat $T > S$ of dimension greater by 1, define the hypersurface $h(S,T)$ to be the union over $\mathbf{x} \in S$ of the cones $K_T(\mathbf{x})$. Then the set non is the cone in $(\mathbb{R}^+)^d$ (equivalently the subset of the positive orthant in $\mathbb{RP}^{d-1}$) consisting of all surfaces $h(S,T)$. In particular, $\mathbf{r} \in \mathsf{non}$ whenever $\mathbf{r}$ is in the relative boundary in $S(\sigma)^\perp$ of K_σ for some $\sigma \in \Sigma$.*

Proof: Given $\mathbf{x} \in V_S$ with $S < T$ as above, suppose $\mathbf{r}$ is in the intersection of $(\mathbb{R}^d)^+$ with the linear span of $\{\tilde{\mathbf{b}}_j : j \in T\}$. For that $\mathbf{r}$, the critical point of the height function $f_\mathbf{r}$ on the flat V_T is then in the linear span of $\{\tilde{\mathbf{b}}_j(\mathbf{x}) : j \in T\}$, which is the condition for $\mathbf{x}$ to be the saddle $\sigma(T)$. Thus for this $\mathbf{r}$ we see that $\sigma(T) \in S$ and hence $\mathbf{r} \in \mathsf{non}$. This proves that $\bigcup h(S,T) \subseteq \mathsf{non}$.

Conversely, if $\mathbf{r} \in \mathsf{non}$ then pick $S < T$ with $\sigma(T) \in S$ and let $\mathbf{x} = \sigma(T)$. Then at that $\mathbf{x}$ we find $\mathbf{r}$ to be in the linear span of $K_T(\mathbf{x})$, proving the reverse inclusion.

Finally, to see that each $h(S,T)$ is a hypersurface, we describe these sets as follows. Denote the ambient space as V, and consider the bilinear mapping $\Delta : V \times V^* \to \mathbb{R}^d$ given by

$$((e_1, \ldots, e_d), (f_1, \ldots, f_d)) \mapsto (e_1 f_1, \ldots, e_d f_d)$$

(the basis in V is fixed). Given the pair (S,T) we let $C_+^* \subseteq V^*$ denote the space of linear functionals constant on V_T and let $L_- \subseteq V$ be the linear space spanned by V_S. The image of the restriction of Δ to $L_- \times C_+^*$ is an immersed hypersurface in $\mathbb{R}^d$, as the rank of the Jacobian of this restriction is $(d-1)$.

Remark: If the dimension of V_- or codimension of V_+ is one, then the image hypersurface is in fact a hyperplane. This must always be the case in dimensions up to three. In general the image surface is not a hyperplane, being a restriction of a quadratic mapping.

The decomposition of gen into components is in general a finer decomposition than is necessary: there may be facets of non across which no discontinuities occur in any terms of the asymptotic formula for $a_\mathbf{r}$. While Proposition 6.1 gives some understanding of the structure of gen, the problem of effectively and efficiently listing chambers is by no means trivial and has been studied in various contexts. Even counting the number of chambers is cited in [DS03] as an open problem, stated in [Ki01, page 57]. That problem, it should be noted, concerned the chamber complex for Kostant's partition function, which is from our point of view simple in that all hyperplanes of $\mathcal{A}$ pass through $(1, \ldots, 1)$. Hence $\tilde{\mathbf{b}}_j \equiv \mathbf{b}_j$ for all j and $\mathbf{x}$, and the problem reduces to describing chambers of a dual hyperplane arrangement; the geometry of non may be more complicated if not all the hyperplanes pass through a single point.

6.2 Finding the highest saddle in contrib

The first three steps of Algorithm 1.1 involve finding all the flats of $\mathcal{A}$, finding the saddle in each flat, and sorting these by according to their height, $f_{\mathbf{r}}(\sigma)$. Here, we consider briefly the complexity of these computations.

One might imagine doing these computations in two ways: for a given $\mathbf{r}$, or simultaneously for all $\mathbf{r}$. The asymptotic formulae we have derived are uniform estimates over compact subsets of the chambers of gen. Typically there are discontinuities or non-analyticities across these boundaries, so a formula valid for all $\mathbf{r}$ is nothing more than a description of the chambers of gen together with a compilation of formulae valid in chambers of gen. Having done the best we can to describe the chambers of gen, we assume henceforth that $\mathbf{r}$, or its chamber, is fixed.

Consider now the problem of finding the highest saddle(s) in contrib. Given Σ in some kind of list form, the problem of height ordering the elements of Σ involves testing inequalities among logarithms of solutions to linear equations. If all the b_{ij} are rational, then these are logarithms of rational numbers, so, exponentiating, the problem is reduced to testing inequalities among algebraic numbers. One need not be satisfied with numerical testing here: rigorous testing algorithms are available; see [GS96].

The complexity of this task in general seems no less than the complexity of listing Σ and sorting by means of an efficient sorting algorithm which calls either a numeric or a rigorous testing procedure for each comparison. There is, however, one special case in which the problem of finding the highest saddle in contrib can be shown to be solvable in polynomial time, namely when the coefficients of F are known to be nonnegative.

Let $\mathcal{D}$ be the domain of convergence of the power series for F about the origin, and let $\mathcal{D}_+$ be the intersection of this with the positive orthant in $\mathbb{R}^d$. Then $\partial\mathcal{D}$ is the union of tori, each one of which intersects $\mathcal{D}_+$ in a unique point, $\mathbf{x}$. Since F is meromorphic, there is a pole somewhere on each of such torus, so the power series is not absolutely convergent on these tori. As $\lambda \uparrow 1$, $F(\lambda\mathbf{x}) \to +\infty$ by nonnegativity of the coefficients and divergence of the power series, whence $\mathbf{x} \in \partial\mathcal{D}_+$. We conclude that the interior of the region $B_0 \in \mathcal{B}$ in the positive orthant with the origin in its closure is contained in $\mathcal{D}$. It follows that if $\sigma \in \partial B_0$ for some $\sigma \in \Sigma$, then σ is a (weakly) minimal point in the terminology of [PW01], meaning that it is on the boundary of a polydisk whose interior is in the domain of convergence of F. Consequently, there can be at most one such $\sigma \in \partial B_0$ (though there may be others on the boundary of the same polydisk if the coefficients of F have periodicity).

On the other hand, let σ be the location of the minimum of the convex function $f_{\mathbf{r}}$ on B_0. Then σ is in the relative interior of precisely one face of B_0, and it follows that $\sigma = \sigma(S)$ for the flat S of which this face is a subset with nonempty relative interior. We see that there is precisely one $\sigma \in \Sigma \cap \partial B_0$ and that it is where $f_{\mathbf{r}}$ is minimized on B_0. The problem of minimizing a smooth convex function on a polytope is solvable in polynomial time by means of interior point methods. It is possible that specific polynomial time bounds such as are given in [NN94] may be improved upon due to the special nature of our convex objective function. In particular, we may explicitly compute the mimimum on any face. Algorithms designed to produce solutions within ϵ of optimal at a time cost of polylog(ϵ) may therefore be told at some point to quit and produce the exact minimum. How best to do this remains a topic for further investigation.

6.3 Nonlinear pole sets

Much of our analysis is valid in a more general setting where the pole variety is the union of smooth components but these are no longer required to be flat. In particular, the five steps beginning with equation (1.2) are valid whenever the pole variety is locally bi-analytically equivalent to a hyperplane arrangement. An additional layer of computation is incurred in Lemma 5.3 since computing higher derivatives of forms will involve the higher derivatives of the linearizing diffeomorphism. The leading contribution will, however, be unchanged as long as it the numerator is locally non-vanishing:

Proposition 6.2 *Let H be a product of linear functions vanishing at the point p and not at the origin and let C be a nonzero cycle local to p in the complement of the zero set V_H of H. Let ϕ be an analytic map on a neighborhood of p whose derivative at p is the identity, let η be analytic and nonvanishing near p, define $F = \eta/H$ and $\tilde{F} = F \circ \phi$. Then*

$$\int_C \mathbf{z}^{-\mathbf{r}} F \, d\mathbf{z} \sim \int_C \mathbf{z}^{-\mathbf{r}} \tilde{F} \, d\mathbf{z}$$

as functions of $\mathbf{r}$.

Proof: Changing coordinates by ϕ in the second integral yields $\int_{C'} (\tilde{\eta}/H) \mathbf{z}^{-\mathbf{r}} \, d\mathbf{z}$ for some $\tilde{\eta} = (1 + h)\eta$ and $h(p) = 0$. The cycle C' is homologous to C. We then have that the difference of the integrals is $\int_C (h\eta/H) \mathbf{z}^{-\mathbf{r}} \, d\mathbf{z}$ and we may apply, for instance (3.8), to see that this is $O\left(|\mathbf{r}|^{-1} | \int_C F \mathbf{z}^{-\mathbf{r}} \, d\mathbf{z}\right)$. $\square$

Suppose, for example, that the denominator of F is a toric polynomial, that is, the product of some number k of binomials: $\prod_{j=1}^k (1 - \mathbf{x}^{\mathbf{m}^{(j)}})$. The toric denominator (together with a unimodularity restriction) is precisely the problem considered in [DS03], where it is applied to provide an exact count of nonnegative integer solutions to simultaneous linear equations. De Loera and Sturmfels use various algebraic-geometric methods to find this piecewise polynomial and to do so more rapidly than existing methods for this "benchmark" problem. To solve this problem in our framework, we would need to apply cohomological reduction to the form $\mathbf{z}^{-\mathbf{r}} F d\mathbf{z}/\mathbf{z}$ near the point $\mathbf{1} := (1, \ldots, 1)$. This is algorithmically straightforward, since the linearizing diffeomorphism at $\mathbf{1}$ is just the logarithm (one must take extra care since the domain of convergence is not strictly logarithmically convex). The nonlinear version of Algorithm 1.1 may therefore be carried out for this problem.

At this point our emphasis in on existence of an algorithm for as wide a class as possible, rather than on finding a good algorithm for a particular problem. Let us compare our methods to those used to solve a particular case of unimodular counting. The Birkhoff polytope $\mathcal{B}_n$ is the polytope in $(\mathbb{R}^+)^{n^2}$ of $n \times n$ doubly stochastic matrices. The Ehrhardt polynomial of the Birkhoff polytope is the polynomial $H_n(t)$ counting the number of rational points with denominator t of the Birkhoff polytope. This turns out to be equal to $a_{t\mathbf{1}}$ where $\{a_{\mathbf{r}}\}$ are coefficients of a generating function whose denominator is a toric polynomial. The leading term of $H_n(t)$ is, properly normalized, the volume of the Birkhoff polytope. The most efficient determination to date of this volume is carried out in [BP03] by methods similar to ones in the present paper. Computing this single integer is hard enough that in [BP03], Beck and Pixton can only get up to $n = 9$ (a later note solving the $n = 10$ case with the equivalent of 17 years of 1 GHz computing time was recently web posted). If we are interested only in computing the leading term, Proposition 6.2 allows us to replace the actual generating function by its linearization near $\mathbf{1}$. We may then use Theorem 2.5 to compute the leading term. We do not expect the procedure outlined in Section 5.1 to do any better than that in [BP03]. In fact the discussion after [BP03, Corollary 4] indicates that the cohomological reductions in Section 5.1 mirror the steps of their algorithm before they apply simplifying tricks. On the other hand, we demonstrate an effective procedure for all instances of generating functions with affine pole sets, whereas the methods in [BP03] seem tailored to the Birkhoff polytope problem in a somewhat *ad hoc* way. For the computation of the entire piecewise polynomial whose generating function has poles on a toric variety, we are virtually certain that the nonlinear version of our method does not perform as well as those in [DS03].

Finally, we turn to the most general nonlinear case that could be handled by our methods, namely when the pole variety is a union of smooth hypersurfaces. In general, application of the five steps of this program requires an effective means of accomplishing the representation of T in the quasi-local basis. At present, we know how to do this only for specific cases such as hyperplane arrangements, toric varieties and smooth algebraic curves [HP02]. We hope to make further progress on this in the near future.

6.4 Non-generic directions

It is somewhat unsatisfying that our formulae are restricted to the interiors of chambers and can become invalid near non. For asymptotics as $\mathbf{r} \to \infty$ in non, the only known results are given in [PW04]. There the singularity, already in normal form, is resolved completely, and the region of integration becomes the product of an open disk with a simplex. When $\mathbf{r} \in \mathsf{non}$, the stationary phase point is located on the boundary of the simplex and asymptotics are available via halfspace and other more complicated stationary phase integrals.

In short, for directions in non, there are results available which are technically difficult and are restricted to the case where the highest $\sigma \in \mathsf{contrib}$ is a minimal point. Unification of these results with the formulae on the interiors of chambers is even more problematic. One would like to find a scaling exponent α so that if the distance from $\mathbf{r}$ to non is or order $|\mathbf{r}|^\alpha$, then one has a formula for $a_\mathbf{r}$ that smoothly interpolates between the behavior in the directions in non and behavior in directions in the interior of the nearby chamber. The first results we know of in this regard were obtained in [BFSS01], where a Airy-type scaling limit was derived. General results in this direction were then obtained by Lladser [Ll04]. Aside from these results, nothing appears to be known, and the area cries out for further research.

7 Appendix

7.1 Relative homology

We begin by quoting the basic lemma of Stratified Morse Theory [GM88, Theorem I:3.2].

Proposition 7.1 (no topological change between critical values) *If $[a, b]$ contains no critical value of the Morse function f, then $\mathcal{M}_a$ is a strong deformation retract of $\mathcal{M}_b$.* □

Recall we have denoted $\xi = \mathbf{z}^{-\mathbf{r}-\mathbf{1}} F(\mathbf{z}) d\mathbf{z}$. The previous proposition then yields:

Corollary 7.2 *If $b < \mathsf{low}$ and $C \in H_d(\mathcal{M}_b)$ then $\int_C \xi = 0$ for sufficiently large $\mathbf{r}$.*

Proof: For any $a < b < \mathsf{low}$, the manifold $\mathcal{M}_a$ is a strong deformation retract of $\mathcal{M}_b$. Thus any integrating j-cycle C in $\mathcal{M}_b$ can be retracted into any lower $\mathcal{M}_a$. This can in fact be done so as to result in a chain C_a not coming too close to the poles of F and possessing not too much volume: the exact statement is that

$$\int_{C_a \cap A(a,2a)} F(\mathbf{z}) d\mathbf{z} < P(a)$$

where $A(a, 2a)$ is the set of $\{\mathbf{z} : a < |\mathbf{z}| < 2a\}$ and P is a polynomial. When $\mathbf{r}$ is sufficiently large, we then see that

$$\int_{C_a} \mathbf{z}^{-\mathbf{r}} F(\mathbf{z})\, d\mathbf{z} \to 0 \tag{7.1}$$

as $a \to \infty$; since the integral is independent of a, it is zero. □

Recalling that $(\mathcal{M}, -\infty)$ is the inverse limit of the homotopy equivalent space $(\mathcal{M}, \mathcal{M}_c)$ for $c < \mathsf{low}$, we may now formalize the fact that cycles in $\mathcal{M}_b$ may be ignored.

Proof of 4.1: Suppose C_1 and C_2 are chains with $[C_1] = [C_2]$ in $H_d(\mathcal{M}, -\infty)$. By hypothesis there are chains D and E_b on $\mathcal{M}$ and $\mathcal{M}_b$ respectively such that $\partial D = C_1 - C_2 - E_b$. By Stokes' Theorem, the integral of any holomorphic form over ∂D vanishes, so we see that

$$\int_{C_1} \mathbf{z}^{-\mathbf{r}} \omega - \int_{C_2} \mathbf{z}^{-\mathbf{r}} \omega = \int_{E_b} \mathbf{z}^{-\mathbf{r}} \omega$$

whence the first result follows from the fact that $f \leq |\mathbf{r}| f_{\hat{\mathbf{r}}}(b)$ on $\mathcal{M}_b$, with $f_{\hat{\mathbf{r}}}(b) < \mathsf{low}$.

For part (ii), let C be a chain in $\mathcal{M}$ which is a cycle in $(\mathcal{M}, -\infty)$ homologous to 0. by the argument in part (i), $C = \partial D + E_b$ where $E_b \in \mathcal{M}_b$ for any b. Since $\int_D \omega = 0$ and $\int_{E_b} \omega \to 0$ as $b \to -\infty$, we see that $\int_C \omega = 0$. $\square$

7.2 Proof of the quasi-local decomposition lemma

A standard Morse theoretic argument is that if f is a smooth height function with finitely many critical points, all hyperbolic, then an arbitrarily small perturbation f_ϵ of f may be chosen so that all the critical values of f_ϵ are distinct and such that the family $\{f_t : 0 \leq t \leq \epsilon\}$ induces an isomorphism on all the subsequent Morse theoretic constructions. For details, see for example the critical point exchange construction in [Vas01, Section 12.4]. We may therefore assume without loss of generality that the critical values of f are distinct. The topology of $(\mathcal{M}, -\infty)$ may now be constructed inductively, up to homotopy equivalence, by observing the change in $\mathcal{M}_c$ as c increases past a critical value.

Proposition 7.3 (cycles across critical values generate) *Suppose that the critical values of f are distinct and let S be the set of critical points of f. For each $x \in S$, then choose $a_x < f(x) < b_x$ for which the interval $[a_x, b_x]$ contains no other critical value of f. Let π_x denote the map on homology induced by the projection $(\mathcal{M}_{b_x}, -\infty) \to \mathcal{M}_{b_x}, \mathcal{M}_{a_x})$. Then there are cycles $\{C_{x,j} : x \in S, j \leq n(x)\}$ such that for each x the cycles $\{\pi_x(C_{x,j}) : j \leq n(x)\}$ are independent in $H_i(\mathcal{M}_b, \mathcal{M}_a)$, and such that $\{C_{x,j} : x \in S, j \leq n(x)\}$ generate $H_i(\mathcal{M}, -\infty)$.*

Proof: We show inductively on c that if c is not a critical value, then we may choose $\{C_{x,j} : f(x) \leq c, j \leq n_x\}$ generate $H_i(\mathcal{M}_c, -\infty)$. This is true if $c < \mathsf{low}$ since then there is nothing to generate. Proposition 7.1 shows that if it is true for a, it is true for any b such that $[a, b]$ contains no critical value of f. Thus it suffices inductively to assume it for $a = a_x$ and prove it for $b = b_x$.

Accordingly, we suppose is true $a = a_x = f(x) - \epsilon$ and set $b = b_x = f(x) + \epsilon$. Pick $\gamma \in H_i(\mathcal{M}_b, -\infty)$. The short exact sequence

$$0 \longrightarrow (\mathcal{M}_a, -\infty) \longrightarrow (\mathcal{M}_b, -\infty) \longrightarrow (\mathcal{M}_b, \mathcal{M}_a) \longrightarrow 0$$

gives rise to a long exact sequence in which one finds

$$H_i(\mathcal{M}_a, -\infty) \xrightarrow{\iota} H_i(\mathcal{M}_b, -\infty) \xrightarrow{\pi_x} H_i(\mathcal{M}_b, \mathcal{M}_a) . \tag{7.2}$$

Choose a preimage $\{C_{x,j} : j \leq n(x)\}$ of a basis for the image of π_x. These are elements of $H_i(\mathcal{M}_b, -\infty)$. Write $\pi\gamma = \sum_j q_j \pi_x(C_{x,j})$. Then $\gamma - \sum_j q_j C_{x,j}$ is in the kernel of π_x, and hence $\iota\gamma' = \gamma - \sum q_j C_{x,j}$ for some γ'. By induction, γ' is in the span of $\{C_{x',j} : f(x') < f(x)\}$. Thus γ is in the span of $\{C_{x',j} : f(x') \leq f(x)\}$, completing the inductive proof of the proposition. $\square$

It is then shown [GM88, Theorem I:3.5.4] that for each critical point σ and for a and b chosen so that $f(\sigma)$ is the only critical value of f in $[a, b]$, there is a neighborhood $\mathcal{N}$ of σ which may be taken arbitrarily small, so that the inclusion $(\mathcal{M}_b \cap \mathcal{N}, \mathcal{M}_a \cap \mathcal{N}) \to (\mathcal{M}_b, \mathcal{M}_a)$ induces a homotopy equivalence. We denote the d-dimensional homology of either of these spaces by $H_{d,\sigma}(\mathcal{M})$, which is the notation that appears in the quasi-local decomposition lemma.

Next it is shown [GM88, Theorem I:10.7] that the pair $(\mathcal{M}_b \cap \mathcal{N}, \mathcal{M}_a \cap \mathcal{N})$ is homotopy equivalent to a space gotten from $\mathcal{M}_a$ by attaching a CW-complex which is a topological product of pairs

$$(A, B) \times (C, D) \cong (D^\lambda, \partial D^\lambda) \times (l^+_{\mathcal{M}}, \partial l^+_{\mathcal{M}}) , \tag{7.3}$$

where the topological product means

$$(A, B) \times (C, D) = (A \times C, A \times D \cup B \times C) .$$

When computing the homology of an attachement relative to the space before the attachement, the attaching map is irrelevant, so the "Morse data" A, B, C and D suffice to determine $H_*(\mathcal{M}_b, \mathcal{M}_a)$.

In (7.3), the disk A is a small neighborhood of the the critical point in the stratum and B is the part of this where the Morse function f has diminished in value by ϵ. It is then not hard to see that the "tangential Morse data"' (A, B) is always the product of a disk of dimension $d^* - \lambda$ with the pair $(D^\lambda, \partial D^\lambda)$, where $d^* = 2 \dim S$ is the real dimension of the stratum and λ is the index of the critical point, which is always $d = (1/2)d^*$ when f is a locally analytic function. The space C in (7.3) is a neighborhood of the origin in $\mathcal{M}^\perp_S$ and the space D is the portion of C where f has diminished by ϵ.

These facts from [GM88] are true for any stratified spaces, the d-dimensional homology of a complement in $\mathbb{C}^d$ of a complex algebraic variety, we can say more. The complement of a complex variety in $\mathbb{C}^d$ always has homology dimension at most d (see, e.g., [Mil63, Theorem 7.1]). Therefore in the long exact sequence for

$$0 \longrightarrow (\mathcal{M}_a, -\infty) \longrightarrow (\mathcal{M}_b, -\infty) \longrightarrow (\mathcal{M}_b, \mathcal{M}_a) \longrightarrow 0$$

all the homology vanishes above dimension d and (7.2) becomes

$$0 \longrightarrow H_d(\mathcal{M}_a, -\infty) \xrightarrow{\iota} H_d(\mathcal{M}_b, -\infty) \xrightarrow{\pi_x} H_d(\mathcal{M}_b, \mathcal{M}_a) .$$

There is no torsion in the top dimension of the space $(\mathcal{M}_b, \mathcal{M}_a)$ (this follows from the natural orientation on a complex variety and its complement). Thus if each π_σ is surjective, then we see inductively that

$$H_d(\mathcal{M}, -\infty) \cong \bigoplus_{\sigma \in \Sigma} H_d(\mathcal{M}_{b_\sigma}, \mathcal{M}_{a_\sigma}) = \bigoplus_{\sigma \in \Sigma} H_{d,\sigma}(\mathcal{M}) \tag{7.4}$$

although we have no natural way yet to represent this sum until we define π_σ.

It is possible to see that for any complement of a complex variety, each π_x is surjective on H_d, but it is easiest to use the special properties of hyperplane arrangements to do so, as these are necessary anyway for the last step in proving the quasi-local decomposition lemma.

The first special fact we use about hyperplane arrangements is that we may choose explicit cycle representatives $\mathsf{cyc}^{\parallel}$ for $H_{\dim S}(V_S)$ which are actual cycles in $(\mathcal{M}, -\infty)$. Thus the map $i_{\mathbf{x}}$ may be defined so as to be an explicit and natural candidate for an isomorphism between $H_{\mathrm{codim}\,\sigma}(\mathcal{M}^{\perp}_{S(\sigma)})$ and what will turn out to be $\pi_\sigma^{-1}(H_{d,\sigma}(\mathcal{M}))$.

The second special fact about hyperplane arrangements is that the homology of (C, D) is actually the local homology of $\mathcal{M}^{\perp}_S$. It is shown in [GM88, Theorem III:3.2] by "fitting together cycle halves" that any relative cycle τ in $H_{\mathrm{codim}\,S}(l^+_{\mathcal{M}}, \partial l^+_{\mathcal{M}})$ is actually a true local (codim S)-dimensional cycle of $\mathcal{M}^{\perp}_S$ The fact that $\mathcal{M}^{\perp}_S$ retracts to a small neighborhood of the origin means that there is a natural map from $H_{\mathrm{codim}\,S}(\mathcal{M}^{\perp}_S)$ to $H_{\mathrm{codim}\,S,loc}(\mathcal{M}^{\perp}_S)$, which may then be projected to $H_{\mathrm{codim}\,S}(C, D)$. Fitting together halves reverses this map and shows that

$$H_{\mathrm{codim}\,S}(C, D) \cong H_{\mathrm{codim}\,S}(\mathcal{M}^{\perp}_S)$$

which is part (i) of the quasi-decomposition lemma.

Also, mapping $H_{d,\sigma}$ first by this natural map and then by the natural map i_σ provides a map from $H_{d,\sigma}(\mathcal{M})$ to $H_d(\mathcal{M}, -\infty)$ which is inverted by π_σ, thus proving part (ii) of the lemma. $\square$

Proof of Lemma 4.5: To see in general that $|\alpha^{-1}(\sigma)|$ computes the dimension of the local homology, use the fact [GM88, Theorem III:3.5] that the d-dimensional local homology of $\mathcal{M}$ and the 0-dimensional local homology of $\mathcal{M} \cap \mathbb{R}^d$ at σ are equal since both are given by the $(d - 1 - \dim_{\mathbb{R}} \sigma)$)-dimensional homology of the order complex above $S(\sigma)$ relative to the complement of the top flat. The 0-dimensional real local homology is just the number of components intersecting the level set $f(\sigma) + \epsilon$ that do not intersect $f(\sigma) - \epsilon$. This is the number of components containing σ in their boundaries, for which f attains a minimum at σ, or in other words, the cardinality of $\alpha^{-1}(\sigma)$. $\square$

7.3 Retraction of the complement of bd to $-\infty$

Lemma 7.4 *For any neighborhood $\mathcal{N}$ of bd in $\mathcal{M}$, The inclusion of $(\mathcal{M}, \mathcal{M} \setminus \mathcal{N})$ into $(\mathcal{M}, -\infty)$ is a homotopy equivalence and therefore induces a homology isomorphism.*

The proof is via construction of a retraction of $\mathcal{M} \setminus \mathcal{N}$ to $-\infty$.

Lemma 7.5 *There is a continuous vector field v on $\mathbb{R}^d \setminus \mathsf{planes}$ such that*

1. $v \equiv 0$ *on* bd*;*
2. *v respects the strata: for any $\mathbf{x} \in V_S$, $v(\mathbf{x}) \in T_{\mathbf{x}}(V_S)$;*
3. *for $\mathbf{x} \notin \mathsf{bd}$, the product in each coordinate, $x_i v_i(\mathbf{x})$, is strictly positive.*

Proof: Let E_λ denote the L^1-sphere $\{\mathbf{x} \in \mathbb{R}^d : \sum_{i=1}^d |x_i| = \lambda\}$. For sufficiently large λ, the set E_λ is disjoint from bd and intersects each unbounded component U of $\mathcal{M} \cap \mathbb{R}^d$. The set $\mathsf{poles} \cup \mathsf{planes}$ subdivides E_λ into convex polytopes. The vertices p_j of these polytopes vary affinely with λ as $x_j + \lambda y_j$. We include enough vertices so that these polytopes subdivide the subdivision by $\mathsf{poles} \cup \mathsf{planes}$, so that no line segment goes through poles without stopping. Fix a sufficiently large λ and triangulate E_λ by simplices in a way that respects the subdivision by $\mathsf{poles} \cup \mathsf{planes}$ and respects the affine homothety that increases λ. Define $v(x_j + \lambda y_j) := y_j$ for all vertices. Within each simplex of the triangulation extend by convex combination. The resulting field w respects that strata since it is forced to be in each stratum at vertices of the triangulation. Also, $x_i v_i(\mathbf{x})$ is always positive, since the ray $\mathbf{x} + \lambda v(\mathbf{x})$ extends without bound avoiding planes. Finally, extend this to $\mathbb{R}^d \setminus \mathsf{bd}$ by letting $v(\rho \mathbf{x}) = \frac{\rho - \lambda_0}{1 - \lambda_0} v(\mathbf{x})$, where $\lambda_0 = \inf\{\rho : \rho x \notin \mathsf{bd}\}$. This automatically yields a limit of 0 at bd, so extending to be zero on bd finishes the construction. $\square$

Corollary 7.6 *The vector field on $\mathbb{C}^d \setminus \mathsf{planes}$ defined by*

$$w(\mathbf{x} + i\mathbf{y}) = v(\mathbf{x}) + i\mathbf{y}$$

extends v to a continuous vector field on $\mathbb{C}^d$ vanishing precisely on bd, respecting the strata, and such that $\nabla f_{\mathbf{r}}(\mathbf{x}) \cdot w(\mathbf{x}) < 0$ for all $\mathbf{x} \notin \mathsf{bd}$ and all $\mathbf{r}$.

Proof: The constraints on the imaginary parts of each strata are linear, not affine, so $v(\mathbf{x}) + i\mathbf{y}$ respects strata if v respects real strata. Continuity, extension and vanishing exactly on bd are immediate. Since $v_i(\mathbf{x})$ has the same sign as x_i, it follows that $\nabla f \cdot v$ is strictly negative on $\mathbb{R}^d \setminus \mathsf{planes}$, and it follows that $\nabla f \cdot w$ is strictly negative on $\mathbb{C}^d \setminus \mathsf{planes}$. $\square$

Proof of Lemma 7.4: Let $\Phi_t(\mathbf{x})$ denote the image of $\mathbf{x}$ at time t under the flow $\mathbf{x}' = w(\mathbf{x})$. Outside of bd this is a homotopy of $\mathcal{M}$ preserving the strata $\mathcal{S}_S, S \in \mathcal{A}$. Given $\mathcal{N}, \mathbf{r}$ and $c < \mathsf{low}$, there is a t sufficiently large such that $\Phi_t[\mathcal{M} \setminus \mathcal{N}] \subseteq \mathcal{M}_c$. $\square$

Acknowledgement: The second author thanks Vic Reiner and Richard Ehrenborg for helpful discussions concerning Theorem 2.1.

References

[BFSS01] Banderier, C., Flajolet, P, Schaeffer, G. and Soria, M. (2001). Random maps, coalescing saddles, singularity analysis and Airy phenomena. *Rand. Struc. Alg.* **19**, 194–246.

[BP03] Beck, M. and Pixton, D. (2003). The Erhardt polynomial of the Birkhoff polytope. *Discrete Comp. Geom.* **30**, 623–637.

[BR83] Bender, E. and Richmond, L.B. (1983). Central and local limit theorems applied to asymptotic enumeration. II. Multivariate generating functions. *J. Combin. Theory Ser. A* **34**, 255–265.

[BM93] Bertozzi, A. and McKenna, J. (1993). Multidimensional residues, generating functions, and their application to queuing networks. *SIAM Rev* **35**, 239–268.

[BH86] Bleistein, N. and Handelsman, R. (1986). *Asymptotic expansions of integrals,* second edition. Dover: Mineola, NY.

[BV99] Brion, M. and Vergne, M. (1999). Arrangements of hyperplanes I: rational functions and the Jeffrey-Kirwan residue. *Ann. Sci. École Norm. Sup.* **32**, 715–741.

[DS03] De Loera, J. and Sturmfels, B. (2003). Algebraic unimodular counting. *Math. Program., Ser. B*, **96**, 183–203.

[GM88] Goresky, M. and MacPherson, R. (1988). *Stratified Morse Theory.* Ergebnisse der Mathematik und ihrer Grenzgebiete 3. Folge Band 14. Springer-Verlag: Berlin.

[GS96] Gourdon, X. and Salvy, B. (1996). Effective asymptotics of linear recurrences with rational coefficients. In: Proceedings of the 5th Conference on Formal Power Series and Algebraic Combinatorics. *Discrete Math.* **153**, 145–163.

[HP02] van der Hoeven, J. and Pemantle, R. (2002). Effective asymptotics of the Taylor coefficients of bivariate rational functions with smooth poles. *In progress.*

[Kar02] Karloff, H. (2002). *Personal communication.*

[Ki01] Kirillov, A. N. (2001). Ubiquity of Kotska polynomials. In: *Physics and Combinatorics*, Proceedings Nagoya 1999, edited by Kirillov, Tsuchiya and Umemura. World Scientific.

[KY96] Kogan, Y. and Yakovlev, A. (1996). Asymptotic analysis for closed multichain queuing networks with bottlenecks. *Queuing. Sys. Th. Appl.* **23**, 235–258.

[Ll04] Lladser, M. (2004). *Asymptotic enumeration via singularity analysis.* Doctoral dissertation, The Ohio State University.

[Mil63] Milnor, J. (1963). *Morse Theory.* Annals of Mathematics Studies, number 51. Princeton University Press: Princeton.

[NN94] Nesterov, Y. and Nemirovsky, J. (1994). Interior-Point Polynomial Algorithms in Convex Programming. SIAM Studies in Applied Mathematics #13. SIAM: Philadelphia.

[OT92] Orlik, P. and Terao, H. (1992). *Arrangements of hyperplanes.* Springer: Berlin.

[Pem00] Pemantle, R. (2000). Generating functions with high order poles are nearly polynomial. In: *Mathematics and Computer Science: Algorithms, trees, combinatorics and probabilities*, 305–321. Birkhauser: Basel.

[PW01] Pemantle, R. and Wilson, M. (2001). Asymptotics of multivariate sequences, part I: smooth points of the singular variety. *J. Comb. Th., series A* **97**, 129–161.

[PW04] Pemantle, R. and Wilson, M. (2004). Asymptotics of multivariate sequences, part II: multiple points of the singular variety. *Preprint.*

[Roc70] Rockafellar, R. T. (1970). *Convex Analysis.* Princeton Mathematical Series no. 28. Princeton University Press: Princeton, NJ.

[Vas01] Vassiliev, V. (2001). *Introduction to topology.* Student Mathematical Library, volume 14. AMS: Providence.

[Won89] Wong, R. (1989). *Asymptotic approximations of integrals.* Academic Press: Boston.