SANDIA REPORT

SAND2009-2011
Unlimited Release
Printed March 2009

Computational Algebraic Geometry for Statistical Modeling FY09Q2 Progress

Philippe Pébay, J. Maurice Rojas, David Thompson

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94-AL85000.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.osti.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online ordering: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2009-2011
Unlimited Release
Printed March 2009

Computational Algebraic Geometry for Statistical Modeling
FY09Q2 Progress

Philippe Pébay
Sandia National Laboratories
M.S. 9159, P.O. Box 969
Livermore, CA 94551, U.S.A.
pppebay@sandia.gov

J. Maurice Rojas
Sandia National Laboratories
M.S. 9159, P.O. Box 969
Livermore, CA 94551, U.S.A.
jmrojas@sandia.gov

David Thompson
Sandia National Laboratories
M.S. 9159, P.O. Box 969
Livermore, CA 94551, U.S.A.
dcthomp@sandia.gov

Abstract

This is a progress report on polynomial system solving for statistical modeling.

Contents

1 Introduction 7
1.1 Motivation 7
2 Polynomial Solvers 11
2.1 From Polynomials to Polyhedra and Back Again 11
2.2 $\mathscr{A}$-Discriminants and Chamber Cones 14
2.3 Implementation of the Chamber Cone Method 21
3 Generating Statistical Models 23
3.1 Drawing samples 23
3.2 Distributions of interest 25
3.3 Modeling Shock Response Spectra 26
References 30

This page intentionally left blank

1 Introduction

This is a progress report on polynomial system solving for statistical modeling. This quarter we have developed our first model of shock response data and an algorithm for identifying the chamber cone containing a polynomial system in n variables with $n+k$ terms within polynomial time – a significant improvement over previous algorithms, all having exponential worst-case complexity. We have implemented and verified the chamber cone algorithm for $n+3$ and are working to extend the implementation to handle arbitrary k. Later sections of this report explain chamber cones in more detail; the next section provides an overview of the project and how the current progress fits into it.

1.1 Motivation

Algebraic Statistics is the art of framing discrete statistical models as polynomial systems (see, e.g., [PW01, PS05, DS08]). These systems tend be sparse, multivariate, and of high degree. Statistical software packages generally solve these systems using Gröbner bases or homotopy continuation — techniques that generate all possible solutions to the system, both real and complex-valued — but run impractically slowly without artistry depending subtly on the underlying problem. Alternatively, optimization techniques are sometimes used, but the known non-algebraic techniques are not guaranteed to find all possible solutions, thus leaving the possibility that the model is not properly fit. We aim to improve the situation by developing a technique that computes all *real*-valued roots to these systems and does so (provably) quickly.

In order to understand the issues more clearly, we will now discuss how these polynomial systems arise and how their structure may be exploited. The polynomial systems we are interested in are generated by considering the probability of observing experimental (simulated or real-world) data.

Example 1.1. Consider an experiment is run several times under varying conditions with the following results

$$O = (a,b,b,c,d,a,d,a).$$

The observations in O consist of four possible *outcomes*: a, b, c, and d. In a given list of observations, each outcome has a multiplicity – the number of times it appears in O. Assuming that the observations are realizations of a *discrete* random variable, the likelihood L of obtaining *all* of those results is the product of the probability of each occurring:

$$L(O) = p(a)p(b)p(b)p(c)p(d)p(a)p(d)p(a),$$

where p denotes the probability function of the random variable. If one is willing to assume that the order of the observations does not matter (i.e., the experimental process is not dependent on the order in which experiments are run), then L can be written

$$L(O) = p(a)^3 p(b)^2 p(c)^1 p(d)^2.$$

Now take the log to get the log-likelihood ℓ:

$$\ell(O) = \log L(O) = 3\log(p(a)) + 2\log(p(b)) + 1\log(p(c)) + 2\log(p(d)).$$

If each probability $p(a)$, $p(b)$, ... is a function of some unknown parameters x_i, then the parameters that maximize ℓ (and thus L) may be determined by taking the derivative of ℓ with respect to the parameters x_i and setting it to the 0 vector. The result[1]

$$\frac{\partial \ell}{\partial x_i}(O) = \frac{3}{p(a)}\frac{\partial p}{\partial x_i}(a) + \frac{2}{p(b)}\frac{\partial p}{\partial x_i}(b) + \frac{1}{p(c)}\frac{\partial p}{\partial x_i}(c) + \frac{2}{p(d)}\frac{\partial p}{\partial x_i}(d) = 0 \tag{1.1}$$

is a set of rational functions each of which can be written as the ratio of two polynomials $n(x_i)/d(x_i) = 0$. If we find x_i^* where all the $n(x_i^*) = 0$, then we have either the most likely set of parameters x_i^* (assuming the Hessian is positive definite), the least likely set of parameters (where the Hessian is negative definite), or a degenerate critical point (when the Hessian is semidefinite)..

In order to solve such polynomial systems, as obtained by following the same methodology as in the above example, Pachter and Sturmfels [PS05] (and others) use Gröbner basis decompositions of the polynomials. This becomes computationally intractable quickly as the degree of the system increases because all of the possible solutions to the system (both real and complex) must be computed. Other approaches do not attempt to find all the roots of the system but, given a starting point, use gradient-descent techniques to find a *single* real root – which may or may not correspond to the global minimum.

The purpose of this project is to develop techniques for solving polynomial systems with complexity depending mainly on the number of *real* roots via a clever use of $\mathscr{A}$-discriminant theory, and to guarantee that *all* the real roots in some given parameter space are identified. It turns out that discriminants allow us to partition the space of unknowns (x_i) into regions with similar behavior.

Example 1.2. A familiar example is the simple univariate, quadratic trinomial

$$f(x) = ax^2 + bx + c,$$

whose discriminant is $\Delta = b^2 - 4ac$. If $\Delta < 0$ then $f(x)$ has 2 complex roots. Where $\Delta > 0$, there are 2 real roots. And when $\Delta = 0$, there is a double real root. $f(x)$ can be rewritten as follows in order to remove one coefficient:

$$g(x) = x^2 + Bx + C$$

where $B = b/a$ and $C = c/a$ and so the discriminant becomes

$$\Delta = B^2 - 4C$$

The plot of $\Delta(B,C) = 0$ is a 2-dimensional curve[2] that separates the (B,C) plane into 3 regions $(\Delta < 0, \Delta > 0, \Delta = 0)$. Each region is called a discriminant chamber. All the polynomials in the same chamber have the same number of real roots.

While the number of chambers can grow exponentially in the number of monomial terms, we have developed an algorithm that can find a polyhedral approximation to the chamber (a *chamber cone*)

[1] Note that evaluating $\log p$ before differentiating instead of afterwards (as shown) is what makes log-likelihood a useful technique. Don't take this equation too literally.

[2] specifically, a parabola.

containing a given system in polynomial time. Our algorithm is already implemented for a broad class of systems of n polynomials in n unknowns, and as of this writing a completely general implementation is being completed.

Knowing the chamber cone that contains a system $F := (f_1(x_1,\ldots,x_n),\ldots,f_n(x_1,\ldots,x_n))$ is important because the vertex of the chamber cone can be used to construct a family of new systems $H_t := (h_{1,t}(x_1,\ldots,x_n),\ldots,h_{n,t}(x_1,\ldots,x_n))$ with the following properties:

1. $H_0 := F$
2. the real roots of $G := H_1$ are easy to find and are the same in number as the real roots of our original system F.
3. $H_t(x)$ is a differentiable function on $[0,1] \times \mathbb{R}^n$.

A technique called homotopy can then be used to find the roots of F. More specifically, homotopy continuation deforms the roots of a system we know (G) to the roots of a system (F) we are trying to understand, and does so numerically by a simple application of Newton-Raphson iteration. **However, what distinguishes the homotopies we use is that we specifically construct H so that G and F lie in the same discriminant chamber, with high probability.**

Existing homotopy techniques [SW05] are robust but are also slow because they choose H without paying any attention whatsoever to discriminant chambers. Without knowing if F and G lie in the same chamber, homotopy methods must then track every complex root (including all non-real roots) of H_t from $t = 0$ to $t = 1$. Sparse polynomials tend to have few real roots, regardless of how many complex roots there are. So, for earlier homotopy methods, much of the cost of tracking roots is spent performing work that will be of no use in the end.

By choosing F to lie in the same chamber cone as G, all of the work tracking non-real roots can be eliminated. Our technique is, by construction, probabilistic: our use of chamber cones implies that there will be a small fraction of inputs which do not actually lie in a putative corresponding discriminant chamber. Earlier algorithms could not attain polynomial time for even a provably small fraction of inputs. Moreover, since the problem of counting the number of real roots is **#P**-hard [BRS08], a small set of hard inputs is very likely unavoidable. Initial experiments indicate that our approach is quite practical. Furthermore, with some additional work we may still get useful results even when F does not lie in the discriminant chamber we expect.

The formulation of this method, and obtaining rigorous estimates on success probability, are the research part of the proposal. But in order to answer these questions, we need a source of polynomials generated by physical systems; it is easy to test with random distributions of polynomials but that does not say anything about how useful the technique will be in practice. That is where ASC simulation data will come in. By modeling phenomena such as properties of shock response spectra (SRS) we will obtain systems of equations representative of physical systems. If they tend to lie away from the vertices of chamber cones, then the technique will provide good start systems – unlike other homotopy methods – so that we need only track the real roots of F.

The remainder of this report covers progress on the theory – particularly on homotopy methods – and on the application of the theory to physical models – particularly shock response data.

2 Polynomial Solvers

The chamber cone technique described in the introduction is a new technique that enables homotopy methods to preserve the underlying real topology. (The existence of such methods was a formerly open problem.) In particular, it provides the first algorithm for finding real roots of multivariate systems of polynomial equations that is polynomial-time (in fixed dimension) with high probability.

This section details

- notation and definitions concerning discriminants of polynomial systems (§2.1),
- the implementation of the chamber cone method (§2.3), and
- the foundations of the theory of the chamber cone method (§2.2.2).

Also of note is that the chamber cone provides a potentially new solution to Smale's 17th Problem [Sma00], and this is outlined in Section 2.2.2.

In summary, the primary effort so far has revolved around settling the most difficult theoretical portions of the chamber cone method and writing preliminary implementations. This has eliminated a potential bottleneck for completion, so the project has succeeded in remaining on schedule and providing a radically new and useful algorithm for algebraic equation solving.

2.1 From Polynomials to Polyhedra and Back Again

In order to understand chamber cones in more detail, we will require a thorough discussion of the discriminant of a polynomial system. First, let us review how polyhedra are intimately (and inextricably) related to polynomials.

2.1.1 Amoeba and Lower Binomials

Perhaps the most direct evidence that polynomials are related to polyhedra is the **Archimedean Amoeba Theorem** [GKZ94].

Example 2.1. Consider $f(x,y) := 1 - x^2 + y - x^7y^5 + x^6y^7$. Then defining

$$\text{Amoeba}(f) := \{(\text{Log}|x|, \text{Log}|y|) \mid x,y \in \mathbb{C} \setminus \{0\}\ ,\ f(x,y) = 0\},$$

a short `Matlab` program reveals that Amoeba(f) actually looks like an amoeba, and its tentacles are precisely predicted by the exponents of f:

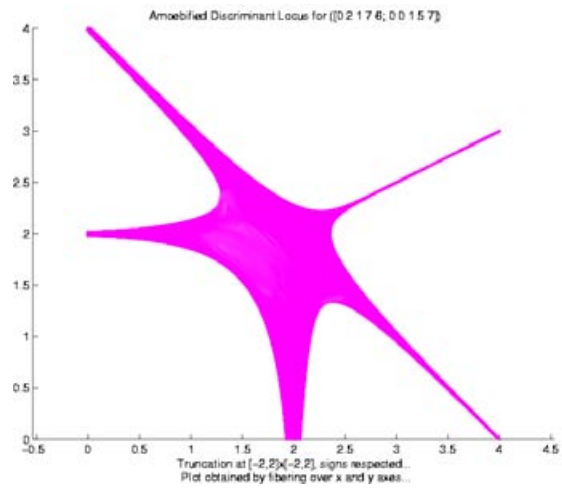

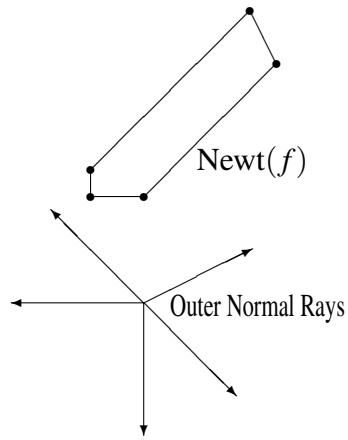

In particular, if we define the **Newton polygon** of f, Newt(f), to to be the convex hull[3] of the exponent vectors of f, we see that the normal vectors of the edges of Newt(f) emanating away from Newt(f) are precisely in the direction of the tentacles! $\diamond$

Amoebae and Newton polygons can easily be extended to arbitrary dimension, and are a basic ingredient in our study. Part of this project advances precise theoretical statements on how the above correspondence can be used to numerically approximate roots of polynomial systems. The most direct evidence of a polyhedral approach to numerical solving can be found already in 1 variable.

Example 2.2. Suppose $f(x) := 1 - cx^{10} + x^{39}$, with c positive. Clearly then, the graph of f can intersect the x axis in one of exactly three ways, depending precisely on the size of c:

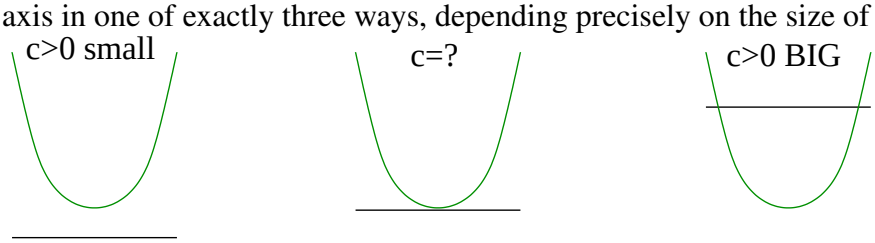

In particular, we see that the existence of a degenerate root is what separates f from having no positive roots, or two positive roots.

More than coincidentally, there is a polyhedral point of view: Let ArchNewt(f) — the **Archimedean Newton polygon** of f — denote the convex hull of the points $(0, -\log|1|)$, $(10, -\log|c|)$, and $(39, -\log|1|)$. It is then easily checked that ArchNewt(f) will either be a line segment or a triangle, and the latter situation falls into two cases, again depending on the size of c:

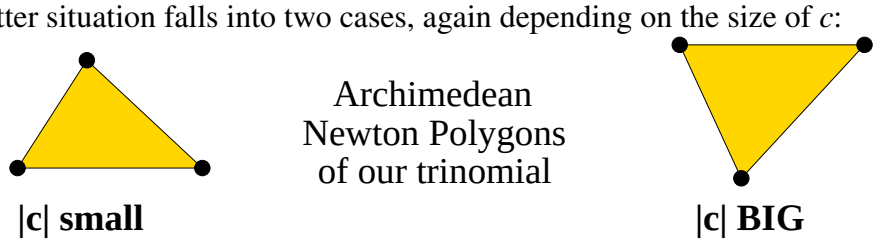

Note also that the **lower hull** of ArchNewt(f) — that is, the union of edges possessing an inner normal with positive last coordinate — naturally defines a subdivision of Newt(f):

[3] i.e., smallest convex set containing...

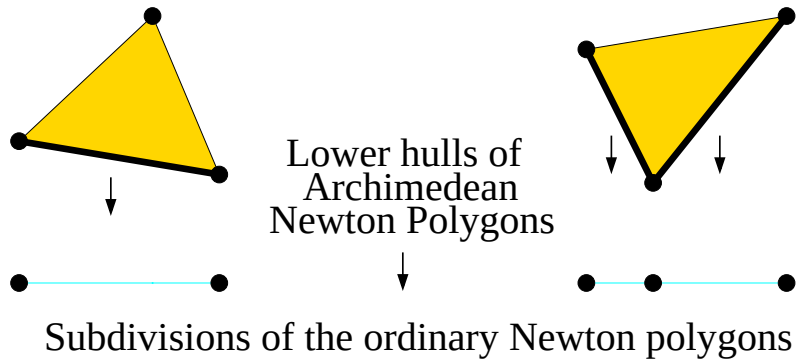

That this has numerical consequences can be seen rather vividly if we fix the magnitude of c and let its phase vary randomly:

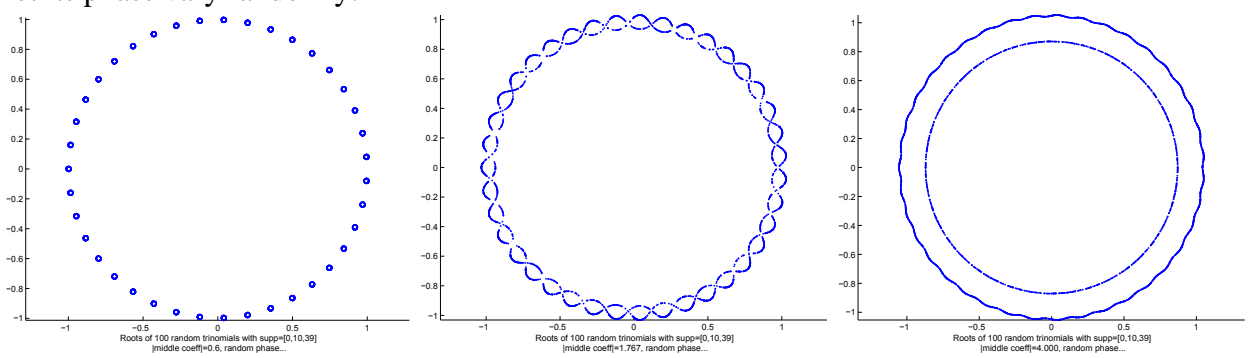

In particular, the key ingredient connecting numerical distribution of roots to our polygonal constructions is the idea of a **lower binomial**: given a subdivision of $\mathrm{Newt}(f)$, we can consider those subsummands of f whose exponents lie in the projection of a lower edge of $\mathrm{ArchNewt}(f)$. For our example, we see that there are two possible sets of binomials for f that can be constructed in this way: $\{1+x^{39}\}$ or $\{1-cx^{10}, -cx^{10}+x^{39}\}$. More to the point, one can easily see that the roots of f appear to cluster around the roots of these binomials! $\diamond$

Constructing simpler equations whose roots somehow approximate the roots of a harder system of equations is exactly our main goal. However, there are many subtleties we must deal with in order to do this rigorously. The first subtlety is that $\mathrm{ArchNewt}(f)$ — as originally defined — does not always yield "lower" binomials whose total number of real roots is the same as the number of real roots of f. However, at least for the trinomial case, this can be corrected easily.

Definition 2.1. For any univariate trinomial $f \in \mathbb{C}[x]$, written $f(x) = a + bx^d + cx^D$ with $0 < d < D$, we define the **normalized** Archimedean Newton polygon of f, $\widetilde{\mathrm{ArchNewt}}(f)$, as the convex hull of $\left(0, \log \frac{D-d}{|a|}\right)$, $\left(d, \log \frac{D}{|b|}\right)$, $\left(D, \log \frac{d}{|c|}\right)$. Finally, when $\widetilde{\mathrm{ArchNewt}}(f)$ is **not** a line segment, we define the **lower binomials** of f to be the subsummands of f whose exponents lie in the projection of a lower edge of $\widetilde{\mathrm{ArchNewt}}(f)$. $\diamond$

Theorem 2.1. *For any real univariate trinomial f, the number of positive (resp. negative) roots of the lower binomials of f is exactly the number of positive (resp. negative) roots of f.* $\square$

This last result is easy to prove, once we have developed sufficient background. Proving a **precise**, analogous statement for numerical root clustering, in complete generality, is more challenging and remains an important focus of this project. However, we have already built (and implemented) algorithms based on higher-dimensional extensions of Theorem 2.1 which we can already prove. To state versions that work for arbitrary systems, however, we will need some more background.

2.2 $\mathscr{A}$-Discriminants and Chamber Cones

We are essentially developing a vast generalization of the following simple fact mentioned in Example 1.2: the polynomial ax^2+bx+c has a 0, 1, or 2 positive roots according as its *discriminant* $\Delta = b^2-4ac$ is negative, zero, or positive. Our generalization of discriminants directly incorporates sparsity[4] to maximize efficiency and, aside from an occasional reference to N-dimensional complex (resp. real) projective space $\mathbb{P}^N_{\mathbb{C}}$ (resp. $\mathbb{P}^N_{\mathbf{R}}$), our development will be decidedly concrete and algorithmic.

The main algebraic geometric object underlying our algorithms will be the *$\mathscr{A}$-discriminant* [GKZ94]. We will see how the $\mathscr{A}$-discriminant is central in polynomial system solving and dictates the underlying algorithmic and numerical complexity. Our main new result is a polynomial-time membership algorithm for **signed chamber cones**. We also have certain probabilistic estimates underway that are relevant to Smale's 17th Problem [Sma00]. First, however, let us gently introduce our setting.

2.2.1 What is an $\mathscr{A}$-Discriminant?

Definition 2.2. Consider a general set of exponents $\mathscr{A}=\{a_1,\ldots,a_m\}\subset\mathbb{Z}^n$ of cardinality m, and the corresponding family of (Laurent) polynomials

$$\mathscr{F}_{\mathscr{A}}:=\{c_1x^{a_1}+\cdots+c_mx^{a_m}\mid c_i\in\mathbb{C}\}$$

where the notations $a_i=(a_{1,i},\ldots,a_{n,i})\in\mathbb{Z}^n$ and $x^{a_i}=x_1^{a_1}\cdots x_n^{a_n}$ are understood. When $c_i\neq 0$ for all $i\in\{1,\ldots,m\}$ then we call $\mathscr{A}$ the *support* of $f(x)=\sum_{i=1}^m c_ix^{a_i}$, also using the notation $\mathrm{Supp}(f)$. $\diamond$

Example 2.3. Taking $m=3$, $n=1$, $(a_1,a_2,a_3)=(2,1,0)$, and $(c_1,c_2,c_3)=(a,b,c)$, we have that $\mathscr{F}_{\mathscr{A}}$ is our familiar family of univariate quadratic polynomials of Example 1.2. $\diamond$

Definition 2.3. Let us define $X_{\mathscr{A}}$ — the *(projective) toric variety associated to $\mathscr{A}$* — to be the closure of the parametrized subvariety $\{[t^{a_1}:\cdots:t^{a_m}]\mid t=(t_1,\ldots,t_n)\in(\mathbb{C}^*)^n\}$ of $\mathbb{P}^{m-1}_{\mathbb{C}}$. We then define $\nabla_{\mathscr{A}}$ — the *$\mathscr{A}$-discriminant variety* [GKZ94, Chs. 1 & 9–11] — to be the closure of the set of all $[c_1:\cdots:c_m]\in\mathbb{P}^{m-1}_{\mathbb{C}}$ such that the hyperplane $\{c_1z_1+\cdots+c_mz_m=0\}$ intersects a regular point of $X_{\mathscr{A}}$ with a tangency. When $\nabla_{\mathscr{A}}$ is a hypersurface, we define (up to sign) the *$\mathscr{A}$-discriminant*, $\Delta_{\mathscr{A}}\in\mathbb{Z}[c_1,\ldots,c_m]$, to be the (irreducible) defining polynomial of $\nabla_{\mathscr{A}}$.

Example 2.4. The preceding definition is merely a formalization of a very old idea: the solutions of a polynomial can always be simultaneously viewed as the intersection of a hyperplane with a

[4] in the sense of a relatively small number of monomial terms as compared to the total polynomial degree.

certain (parametrized) algebraic set. In particular, taking $\mathscr{A}=\{0,1,2\}$ again, we see that $X_{\mathscr{A}}$ is simply a quadratic curve in the projective plane. (The triple of *affine* coordinates $(1,t_1,t_1^2)$ defines a curve in $\mathbb{C}^3$ as we let t_1 range over $\mathbb{C}$. However, the triple of *projective* coordinates $[1:t_1:t_1^2]$ stands for an entire *equivalence class* of affine coordinates of the form $(\lambda,\lambda t_1,\lambda t_1^2)$ where λ is any nonzero complex number.) We leave it as an exercise, after computing the tangent space of $X_{\mathscr{A}}$ at $[z_1:z_2:z_3]$, to see that the projective line $\{c_1z_1+c_2z_2+c_3z_3=0\}$ intersects $X_{\mathscr{A}}$ with a tangency iff $c_2^2-4c_1c_3=0$. $\diamond$

In essence, the $\mathscr{A}$-discriminant variety consists of all those $f\in\mathscr{F}_{\mathscr{A}}$ with f possessing a degenerate complex root, possibly "at infinity." The introduction of $X_{\mathscr{A}}$ then gives us a rigorous and combinatorial means of making sense of roots at infinity.

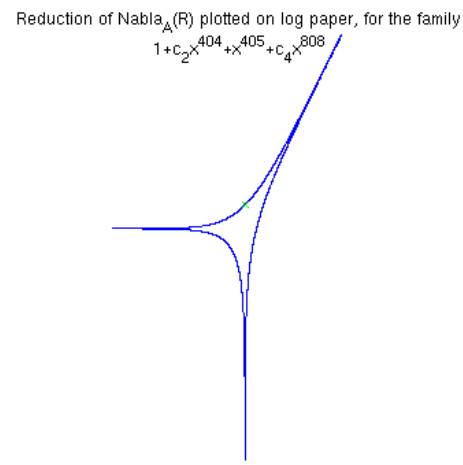

Example 2.5. Taking $\mathscr{A}=\{0,404,405,808\}$, we see that $\mathscr{F}_{\mathscr{A}}$ consists simply of polynomials of the form $f(x):=c_1+c_2x^{404}+c_3x^{405}+c_4x^{808}$. The underlying $\mathscr{A}$-discriminant is then a polynomial in the c_i having 609 monomial terms and degree 1604. In particular, the standard commercial computer algebra system `Maple 12` (running on a dual 2Ghz AMD Opteron system with 16Gb RAM) gives an out of memory error before computing $\Delta_{\mathscr{A}}$. However, while $\Delta_{\mathscr{A}}$ is unwieldy, we can still easily plot the real part of its zero set $\nabla_{\mathscr{A}}(\mathbb{R})$ via the *Horn-Kapranov uniformization* (see Theorem 2.2 below and the illustration to the right). $\diamond$

The plotted curve is then the image of the real roots of $\overline{\Delta}_{\mathscr{A}}:=\Delta_{\mathscr{A}}(1,\gamma_1,1,\gamma_2)$ within Amoeba$(\overline{\Delta}_{\mathscr{A}})$ where, for any polynomial $g\in\mathbb{Z}[c_1,\ldots,c_m]$, we define its *amoeba*, Amoeba(g), to be

$$\{(\log|c_1|,\ldots,\log|c_m|)\ \mid\ g(c_1,\ldots,c_m)=0 \text{ and } c_1\cdots c_m\neq 0\}.$$

Amoebae give us a convenient way to introduce polyhedral methods, such as tropical geometry, into our setting.

$\mathscr{A}$-discriminants are notoriously large in all but a few highly restricted cases. For instance, the polynomial $\overline{\Delta}_{\{0,404,405,808\}}$ defining the curve above has the following coefficient for $\gamma_1^{808}\gamma_2$:
903947086576700909448402875044761267196347419431440828445529608410806270
...[2062 digits omitted]... ...9344158847266670406196231042990817031174921755033 6.
Fortunately, we have the following theorem, describing a one-line parametrization of $\nabla_{\mathscr{A}}$.

Theorem 2.2. *[The Horn-Kapranov Uniformization] (See [DFS07, Prop. 4.1] and [Kap91, PT05].) Given $\mathscr{A}:=\{a_1,\ldots,a_m\}\in\mathbb{Z}^n$, the discriminant locus $\nabla_{\mathscr{A}}$ is exactly the closure of*

$$\left\{[u_1t^{a_1}:\cdots:u_mt^{a_m}]\ \middle|\ u:=(u_1,\ldots,u_m)\in\mathbb{C}^m,\ \mathscr{A}u=\mathbf{O},\ \sum_{i=1}^m u_i=0,\ t=(t_1,\ldots,t_n)\in(\mathbb{C}^*)^n\right\}.\quad\square$$

Thus, once we know the right nullspace of the $(n+1)\times m$ matrix $\hat{\mathscr{A}}$ having $i^{\underline{\text{th}}}$ column the transpose of $(1,a_i)$, we have a formula parametrizing $\nabla_{\mathscr{A}}$.

Example 2.6. Returning to Example 2.3 with the univariate quadratic trinomial $f(x) = ax^2 + bx + c$ where $(a,b,c) \in \mathbb{C}^3$, then with notation of Theorem 2.2, one has $\mathscr{A} = \{0,1,2\}$ and $\hat{\mathscr{A}} = \begin{bmatrix} 1 & 1 & 1 \\ 2 & 1 & 0 \end{bmatrix}$ whose right nullspace is generated by $(1,-2,1)$. Therefore, one sees from Theorem 2.2 that the discriminant variety of f is exactly the closure of $[t^2 : -2t : 1]$ where $t \in \mathbb{C}^*$ or, in other words, the closure of the set of all $(a,b,c) \in \mathbb{C}^3$ such that $b^2 - 4ac = 0$. One thus retrieves the usual definition of the discriminant for the quadratic trinomial. The interested reader can do a similar calculation to retrieve a *parametric* formula for the zero set of the much larger cubic discriminant polynomial. The latter polynomial was derived in the 17th century independently by Tartaglia and Cardano. $\diamond$

Example 2.7. Continuing our Example 2.5, we observe that $\hat{\mathscr{A}} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 0 & 404 & 405 & 808 \end{bmatrix}$ has right nullspace generated by $(1,-405,404,0)$ and $(1,-2,0,1)$. Theorem 2.2 then tells us that $\nabla_{\mathscr{A}}$ is simply the closure of the rational surface

$$\left\{ \left[\lambda_1 + \lambda_2 : -405\lambda_1 - 2\lambda_2 t^{404} : 404\lambda_1 t^{405} : \lambda_2 t^{405}\right] \;\middle|\; \lambda_1, \lambda_2 \in \mathbb{C}, t \in \mathbb{C}^* \right\}$$

in $\mathbb{P}^3_{\mathbb{C}}$. Note that f and $\frac{1}{c_1} f$ have the same roots and that $u \mapsto u^{1/405}$ is a well-defined bijection on $\mathbb{R}$ that preserves sign. Note also that the roots of f and $\bar{f}(y) := \frac{1}{c_1} f\left(\left(\frac{c_1}{c_3}\right)^{1/405} y\right)$ differ only by a scaling when f has real coefficients, and that $\bar{f}$ is of the form $1 + \gamma_1 y^{404} + y^{405} + \gamma_2 y^{808}$. It then becomes clear that we can reduce the study of $\nabla_{\mathscr{A}}(\mathbb{R})$ to a lower-dimensional slice. More precisely, intersecting $\nabla_{\mathscr{A}}$ with the plane defined by $c_1 = c_3 = 1$ yields the following parametrized curve in $\mathbb{C}^2$:

$$\overline{\nabla}_{\mathscr{A}} := \left\{ \left(\frac{-405\lambda_1 - 2\lambda_2}{\lambda_1 + \lambda_2} \left(\frac{404\lambda_1}{\lambda_1+\lambda_2}\right)^{-404/405}, \; \frac{\lambda_2}{\lambda_1+\lambda_2} \left(\frac{404\lambda_1}{\lambda_1+\lambda_2}\right)^{-808/405} \right) \;\middle|\; \lambda_1, \lambda_2 \in \mathbb{C} \right\}.$$

In other words, the preceding curve describes the set of all polynomials of the form $1 + \gamma_1 x^{404} + x^{405} + \gamma_2 x^{808}$ having a degenerate root. Our preceding illustration of the image of $\overline{\nabla}_{\mathscr{A}}(\mathbb{R})$ within $\mathrm{Amoeba}(\overline{\Delta}_{\mathscr{A}})$ thus has the following explicit parametrization:
$\left\{ \left(\log|405\lambda_1 + 2\lambda_2| - \frac{1}{405}\log|\lambda_1 + \lambda_2| - \frac{404}{405}\log|404\lambda_1|, \log|\lambda_2| + \frac{403}{405}\log|\lambda_1 + \lambda_2| - \frac{808}{405}\log|404\lambda_1|\right) \right\}_{[\lambda_1:\lambda_2] \in \mathbb{P}^1_{\mathbb{R}}}$ $\diamond$

2.2.2 Discriminant Chambers and the Cayley Trick

Our primary motivation is polynomial system solving over the real numbers and understanding real zero sets. $\mathscr{A}$-discriminants are crucial to fast numerical solving precisely because (a) the distance of a polynomial to $\nabla_{\mathscr{A}}$ is (under mild assumptions) the reciprocal of the condition number [MR04] and (b) the real part of $\nabla_{\mathscr{A}}$ determines where (in coefficient space) a real zero set changes topology.

Definition 2.4. We call any connected component of the complement of $\nabla_{\mathscr{A}}$ in $\mathbb{P}^{m-1}_{\mathbb{R}} \setminus \{c_1 \cdots c_m = 0\}$ a *(real) discriminant chamber*. Also, if $\overline{\nabla}_{\mathscr{A}}$ is the intersection of $\nabla_{\mathscr{A}}$ with a linear subspace H then we call the intersection of any discriminant chamber with H a *reduced* discriminant chamber. $\diamond$

Under mild assumptions on H, one can completely characterize discriminant chambers via reduced discriminant chambers. In particular, we will pick H via certain specially chosen coordinate subspaces.

Definition 2.5. Given any $\mathscr{A}=\{a_1,\ldots,a_m\}\subset\mathbb{Z}^n$ of cardinality $m\geq n+1$, we call a set of indices $C\subseteq\{2,\ldots,m\}$ an *odd cell* iff the $n\times n$ matrix whose $i^{\underline{\text{th}}}$ column is a_i-a_1 has odd determinant. $\diamond$

In essence, an odd cell allows us to pick a subset of coefficients f that we can homogenize to 1 (by rescaling variables) *without* losing any information about the real roots of the family of polynomials $\mathscr{F}_{\mathscr{A}}$. This allows us to fix a useful slice $\nabla_{\mathscr{A}}\cap H$ to study reduced chambers.

Theorem 2.3. *[BRS08] Suppose C is an odd cell of $\mathscr{A}\subset\mathbb{Z}^n$, $H:=\{[c_1:\cdots:c_m]\mid c_i\in\mathbb{C}$ with $c_j=1\iff j\in\{1\}\cup C\}$, and $\overline{\nabla}_{\mathscr{A}}$ is the image of $\nabla_{\mathscr{A}}\cap H$ under the natural projection forgetting coordinates in $\{1,\ldots,m\}\setminus C$. Then:*

1. *$f\in\mathscr{F}_{\mathscr{A}}\setminus\nabla_{\mathscr{A}}$ implies that there is a reduced $\mathscr{A}$-discriminant chamber $\bar{\mathscr{C}}$ and a polynomial $\bar{f}(x)=\left(\sum_{i\in\{1\}\cup C}x^{a_i}\right)+\sum_{i\in\{2,\ldots,m\}\setminus C}c_ix^{a_i}\in\bar{\mathscr{C}}$ such that the real zero sets of f and $\bar{f}$ are diffeotopic.*
2. *Suppose $\bar{\mathscr{C}}$ is any reduced $\mathscr{A}$-discriminant chamber and $\bar{f},\bar{g}\in\bar{\mathscr{C}}$. Then the real zero sets of $\bar{f}$ and $\bar{g}$ are smooth, and diffeotopic to one another. In particular, if $\mathscr{A}\subset\mathbb{Z}^1$ then $\bar{f}$ and $\bar{g}$ have the same number of positive (resp. negative) roots.* $\square$

Example 2.8. As observed earlier in Example 2.7, it suffices to study polynomials of the form $\bar{f}(x)=1+\gamma_1x^{404}+x^{505}+\gamma_2x^{808}$ to study the real roots of real polynomials in $\mathscr{F}_{\{0,404,405,808\}}$. So Example 2.7 merely carried out the special case $(\mathscr{A},C)=(\{0,404,405,808\},\{3\})$ of Theorem 2.3. $\diamond$

Studying reduced chambers still requires keeping track of signs but our preceding amoeba illustration lost this information. Nevertheless, we can still visualize reduced chambers correctly by a simple modification of the definition of amoeba.

Example 2.9. Let us glue together 4 truncated and reflected copies of $\mathbb{R}^2$ (one for each possible pair of signs $(\pm,\pm)$) and plot $(\pm\log|\gamma_1|,\pm\log|\gamma_2|)$ as (γ_1,γ_2) ranges over the roots of $\overline{\Delta}_{\{0,404,405,808\}}$, with the signs on a plotted point given by $(\text{sign}(\gamma_1),\text{sign}(\gamma_2))$. The illustration to the right shows this gluing, with the dark cross separating the copies of $\mathbb{R}^2$, and the lighter thicker region indicating the deletion of sufficiently large coordinates, i.e., limits at infinity. $\diamond$

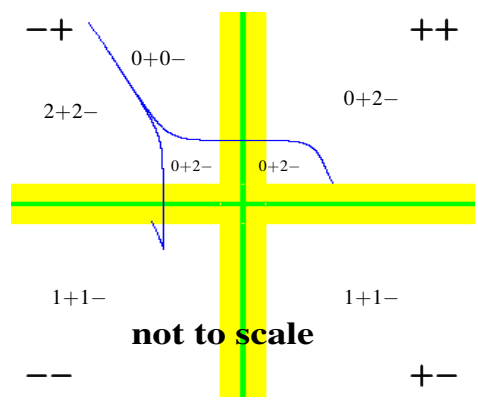

It is worth noting that all reduced chambers in our last illustration are unbounded: 3 merely *appear* to be bounded because of our truncation of large coordinates. We thus see that $\overline{\nabla}_{\mathscr{A}}$ (and the coordinate cross) separates the coefficient space into 8 reduced chambers. More to the point, the number of real roots of $\bar{f}$ is *constant* on any reduced discriminant chamber, as indicated by the numbered labels. (The reduced chamber bounded by the small spike consists of polynomials with exactly 3 positive roots and exactly 1 negative root.) Furthermore, one can prove that any $\bar{f}$ close enough to $\overline{\nabla}_{\mathscr{A}}$ has roots that are poorly separated. For example, the vertex of the lower left spike is

$$(\gamma_1^*,\gamma_2^*)=\left(-\frac{405\sqrt[405]{404\cdot808\cdot403^{404}}}{404^2},-\frac{405\sqrt[405]{404^2\cdot808^2\cdot403^{403}}}{4\cdot404^4}\right)\approx(-1.0166671\ldots,-.0000015832063\ldots)$$

and $\bar{f}^*(x):=1+\gamma_1^*x^{404}+x^{405}+\gamma_2^*x^{808}$ has a root of multiplicity 3 at

$$r := \sqrt[403]{\frac{-405}{404\cdot 808\cdot \gamma_2^*}} \approx 1.0166733769...$$

So, by basic complex analysis, as $\bar{f} \longrightarrow \bar{f}^*$, $\bar{f}$ will have 3 distinct roots converging toward r.

To numerically approximate all the real roots of a polynomial system, one must first know how many there are. The crucial question is then:

Q1: How do we know which chamber we are in, and the number of real roots for the polynomial systems in a given chamber?

Both parts of the question require some care, as our next example illustrates. However, we will first need to introduce the *Cayley Trick*, in order to use $\mathscr{A}$-discriminants to study *systems* of multivariate polynomials.

Definition 2.6. Given any finite points sets $\mathscr{A}_1, \ldots, \mathscr{A}_k \subset \mathbb{Z}^n$, we define their *Cayley embedding*, $\mathrm{Cay}(\mathscr{A}_1, \ldots, \mathscr{A}_k) \subset \mathbb{Z}^{n+k-1}$, to be $(\mathscr{A}_1 \times \{0\}^{k-1}) \cup (\mathscr{A}_2 \times e_{n+1}) \cup \cdots \cup (\mathscr{A}_k \times e_{n+k-1})$ where e_i denote the $i^{\underline{\text{th}}}$ standard basis vector. $\diamond$

The following example (quoted from [DRRS07]) shows that bounded chambers can occur, and extremal examples can reside in exceedingly small bounded chambers. Our example also reveals the existence of codimension ≥ 2 pieces for $\nabla_{\mathscr{A}}(\mathbb{R})$.

Example 2.10. Consider the 2×2 family of system of trinomials

$$F_{a,b} := \begin{cases} f_1(x,y) := x_1^6 + \gamma_1 x_2^3 - x_2 \\ \\ f_2(x,y) := x_2^6 + \gamma_2 x_1^3 - x_1 \end{cases}$$

discovered by Rusek and Shih in [DRRS07]. This system also has a well-defined discriminant, obtainable via the $\mathscr{A}$-discriminant and the Cayley embedding: defining $\mathscr{A}$ to be $\mathrm{Cay}(\mathrm{Supp}(f_1), \mathrm{Supp}(f_2))$, and abusing notation by writing $\mathscr{A}$ as a matrix whenever convenient, we have

$$\mathscr{A} = \begin{bmatrix} 6 & 0 & 0 & 0 & 3 & 1 \\ 0 & 3 & 1 & 6 & 3 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 \end{bmatrix}.$$

Moreover, using the odd cell $C = \{3,4,6\}$, the resulting $\overline{\nabla}_{\mathscr{A}}$ describes exactly those $(a,b) \in \mathbb{C}^2$ where $F_{a,b}$ has a degenerate complex root (possibly at infinity). The image of $\overline{\nabla}_{\mathscr{A}}(\mathbb{R})$ within $\mathrm{Amoeba}(\overline{\Delta}_{\mathscr{A}})$ (save for 3 isolated points), and its unfolded version, are shown below:

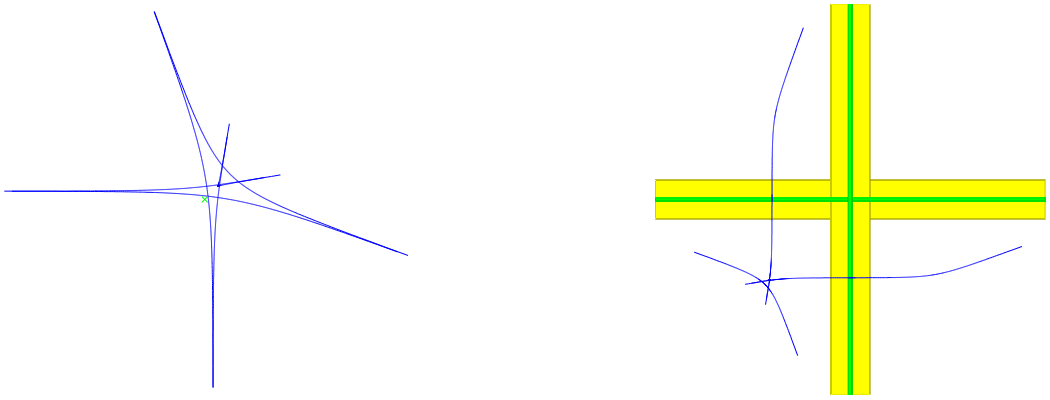

In particular, there is a bounded chamber $\bar{\mathscr{C}}$ (too small to see) in the lower left quadrant that is the *only* chamber containing $F_{a,b}$ having exactly 5 positive roots. (All other $F_{a,b}$ have exactly 3 or 1 positive roots.) Furthermore, $\overline{\nabla}_{\mathscr{A}}(\mathbb{R})$ contains exactly 3 isolated points (not shown) and they lie in the third quadrant [DRRS07]. Note that the isolated points do not affect the connectivity of the reduced chambers. $\diamond$

The tiny chamber in our last example required an extensive computation search. Since it is thus not trivial to determine chamber membership, it is important to approximate chambers by simpler polyhedral objects.

2.2.3 From Chambers to Chamber Cones

Perhaps the simplest observation one can make from our preceding chamber illustrations is that there are large chambers and small chambers, and that "extremal" topological behavior (such as possessing a large number of real roots) appears to occur in extremal chambers. Such a sentiment can in fact be made rigorous.

Definition 2.7. We call a discriminant chamber $\mathscr{C}$ with $\log|\mathscr{C}|$ unbounded an **unbounded chambers**. We also call the smallest polyhedral cone containing $\log|\mathscr{C}|$ a **chamber cone**. $\diamond$

Example 2.11. Returning to our 2×2 system of trinomials from the last section, we see that there are exactly 6 unbounded chambers, shown in the left-hand illustration. (There are also exactly 9 bounded chambers.) The red lines indicate 6 chamber cones. Note in particular how these cones can overlap and **not** share common vertices.

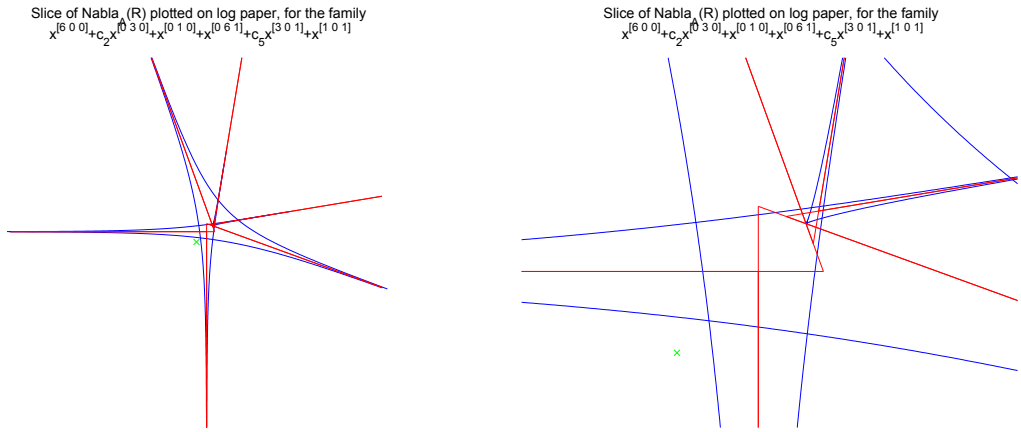

◇

A useful corollary of the Horn-Kapranov Uniformization is the fact that chamber cones always exist and admit simple explicit formulae. This fact could be seen as surprising, considering that $\mathscr{A}$-discriminants are generally intractably large.

Theorem 2.4. *Suppose $\mathscr{A} \subseteq \mathbb{Z}^n$ has cardinality m and let us abuse notation by also considering $\mathscr{A}$ to be an $n \times m$ matrix with $i^{\underline{th}}$ row α_i and $i^{\underline{th}}$ column a_i. Then every unbounded $\mathscr{A}$-discriminant chamber lies in a unique chamber cone. In particular, every chamber cone has facets (called* **walls***), there are exactly m walls, and each wall can be written in the form $\{x \in \mathbb{R}^m \mid N_i(x - s_i) = 0\}$ where $N_i = (N_{1,i}, \ldots, N_{m,i})$ is any generator of the right null space of the $(m-1) \times m$ matrix*

$$\mathscr{W}_i := \begin{bmatrix} 0 & \cdots & 0 & 1 & \cdots & 1 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 \\ 0 & \cdots & 0 & 0 & 1 & \cdots & 1 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & \cdots & 0 & & \ddots & & & \ddots & \ddots & & 0 & \cdots & 0 \\ 0 & \cdots & 0 & 0 & \cdots & 0 & 0 & 1 & \ddots & 1 & 0 & \cdots & 0 \\ & & & & & & \alpha_1 & & & & & & \\ & & & & & & \vdots & & & & & & \\ & & & & & & \alpha_n & & & & & & \\ 0 & \cdots & 0 & 1 & \cdots & & & & \cdots & 1 & 0 & \cdots & 0 \end{bmatrix};$$

there are exactly $m-n-2$ rows preceding the rows containing $\alpha_1, \ldots, \alpha_n$; row j of $\mathscr{W}_i$ has 1s in positions $i+j-1$ through $i+j+m-n-4$ (mod m), and s_i has $j^{\underline{th}}$ coordinate 0 or $\log|N_{j,i}|$ according as $N_{j,i}$ is 0 or not. □

An important consequence of Theorem 2.4 is that while the underlying $\mathscr{A}$-discriminant polynomial $\Delta_{\mathscr{A}}$ may have huge coefficients, the **walls** of $\nabla_{\mathscr{A}}$ admit a concise description involving few bits, with the exception of the transcendental coordinates coming from the "shift" s_i. Another complication with working with chamber cones is that while there are just m walls, the various intersections of walls can easily produce $\Omega(m^n)$ chamber cones in general. Nevertheless, via some care applied to approximations of logarithms and linear forms therein [Ber03, Nes03, BRS08], given a **single** polynomial system, we can quickly find which chamber cone contains it via linear programming.

Theorem 2.5. *Following the notation above, let τ denote the maximum bit size of any coordinate of $\mathscr{A}$, and suppose $f \in \mathscr{F}_{\mathscr{A}}$. Then we can determine the unique chamber cone containing f — or correctly decide if f is contained in 2 or more chamber cones — within a number of arithmetic operations polynomial in m, n, and τ. Furthermore, assuming Baker's refinement of the ABC-Conjecture, we can also attain* **bit**-*complexity polynomial in m, n, and τ.*

At this point, we will not enter into a discussion of the ABC-Conjecture, which would take us too far afield into Diophantine Approximation. (A full discussion of the necessary can be found in [Bak77].) So let us move to the principal result of our project so far.

Theorem 2.6. *Following the notation above, let us write $f(x) = \sum_{i=1}^{m} c_i x^{a_i}$ and define $\mathrm{nanewt}(f)$ to be the convex hull of $\{(a_i, v_i - \log|c_i|) \mid i \in \{1,\ldots,m\}\}$, where $(v_1,\ldots,v_m)$ is the vertex of the unique chamber cone containing f. (We set $(v_1,\ldots,v_m) = (0,\ldots,0)$ should there be not be a unique cone containing f.) Then the real zero set of f has the same topology as the limit of the real zero set of $\hat{f}_t$ as $t \longrightarrow 0$, where $\hat{f}_t(x) := \sum_{i=1}^{m} c_i t^{-v_i} x^{a_i}$. In particular, when f is the Cayley embedding of an $n \times n$ polynomial system F and f lies in a unique chamber cone, then there is a natural collection of lower* **binomial systems** *whose roots approximate the roots of F. Moreover, the number of roots of the resulting collection of lower binomial systems in any orthan $\mathscr{O}$ correctly counts the number of roots of F in $\mathscr{O}$.*

The quality of the numerical approximations obtained via lower binomial systems and chamber cones appears to be excellent in our experiments so far. Nevertheless, we are devoting considerable effort toward finding a precise and provable statement so we can rigorously certify the output of our algorithms. So far, we have already found preliminary estimates for the probability that f lies in a unique chamber cone. Such probabilistic estimates will form an important part of our future deliverables.

Let us now briefly describe how we determine chamber cone membership efficiently.

2.3 Implementation of the Chamber Cone Method

In summary, the chamber cone method has been fully implemented in `Matlab` by Rojas, and will be converted to `C++` in the coming weeks. Further integration of the method into numerical homotopy code, and incorporating lower binomials, will also be done in the coming weeks. We also point out that the chamber cone method has been verified extensively in the special case $m \leq n+3$, and that code for lower binomials has been completed for univariate polynomials and 2×2 polynomial systems with bidegree structure.

So, to conclude our discussion of the chamber cone method, we now detail our Linear Programming (LP) approach. The basic idea is to observe that wherever an $f \in \mathscr{F}_{\mathscr{A}}$ lies in coefficient space, and regardless of which chamber f lies in, f also lies in a unique cell of a **half-plane arrangement**. In particular, due to the structure of the underlying $\mathscr{A}$-discriminant amoeba, the chamber cones are actually determined not by hyperplanes, but by **halves** of hyperplanes. For instance, in our last example, we saw 6 rays determining 6 chamber cones — not 6 lines.

In general, the $i^{\underline{\text{th}}}$ half-plane is determined by N_i, s_i, and an additional datum I_i: the $i^{\underline{\text{th}}}$ **Inequality Vector**. In particular, the $i^{\underline{\text{th}}}$ half-plane consists simply of those points in the $i^{\underline{\text{th}}}$ wall having sufficiently large inner product with I_i. Much like Theorem 2.4, I_i also admits a simple explicit formula.

There is a well-established combinatorial theory of hyperplane arrangements, but the fact that we deal instead with half-planes complicates matters slightly. So our first algorithm has an initial phase that deals with the underlying hyperplane arrangement, and then has a final phase to deal with half-planes. Here is a high-level description:

Input: A subset $\mathscr{A} \subset \mathbb{Z}^n$ and a polynomial $f \in \mathscr{F}_{\mathscr{A}}$. (Note: For polynomial systems, the Cayley Embedding defined earlier in Definition 2.6 gives us an easy way to convert into the format at hand.)
Output: The vertex v of the unique chamber cone containing f (assuming f lies in an unbounded m-dimensional cell of $\mathscr{H}$), or a correct declaration that f does not lie in any unbounded m-dimensional cell of $\mathscr{H}$, where $\mathscr{H}$ is the underlying hyperplane arrangement defined by the walls of $\nabla_{\mathscr{A}}$.

1. For all $i \in \{1, \ldots, m\}$, determine which side of the hyperplane $W_i := \{N_i(x - s_i) = 0\}$ contains f. (If f lies on some W_i then stop and report.) This defines a unique m-dimensional cell C in $\mathscr{H}$ containing f.

2. Observe that C is the feasible set of a collection $\mathscr{S}$ of linear inequalities.

3. **(Removing Misoriented Cell Walls)**
   Using LP unboundedness checks, determine if there is a wall W with its corresponding half-plane **not** contained in C. Remove all inequalities from $\mathscr{S}$ that define such walls.

4. **(Removing Redundant Constraints)**
   Using LP feasibility, remove all inequalities from $\mathscr{S}$ corresponding to walls that do not intersect C.

5. **(Finding Enough Walls to Get a Vertex)** Using LP unboundedness checks, determine if there are $\geq m - n - 1$ (properly oriented) half-planes in the boundary of C. If so, compute their intersection v and stop. Otherwise, report failure.

A preliminary implementation worked quite successfully in the special case $m = n + 3$, and an extension to arbitrary m was completed during the writing of this report. So far, we have successfully experimented with high degree 3×3 systems of trinomials and 4×4 quadratic systems. The speed of the implementation looks promising and we are about to pursue more realistic examples involving 5×5 systems.

3 Generating Statistical Models

In the introduction, we mentioned that a list of observations O is generated with realizations of a random variable. We would like to learn as much as we can about the realizations produced by the random variable. This is typically done by modeling the random variable as one or more probability distributions. In this section, we will discuss techniques for generating and calibrating models; and then we will apply these techniques to some preliminary shock response data. Many of these models have been developed in some form for other MLE techniques – the discussion below is simply meant to show which of these techniques can be used to generate polynomial systems and how to do so.

3.1 Drawing samples

There are two classes of problems we will address and the distinction between them is the form that the log-likelihood $\ell = \log L$ takes on.

3.1.1 Continuous-parameter inference

The first class of problem places 2 requirements on ℓ:

- it must generally be a product of terms without any addition (because $\log(a+b)$ may not be simplified without resulting in terms that do not yield a polynomial system), and
- the parameters of interest must be continuous variables (because we will be differentiating ℓ with respect to each parameter).

The general approach is to simplify ℓ using logarithmic identities and then differentiate ℓ with respect to the parameters of interest.

Example 3.1. An example of this class of problem is approximating O as realizations from a normal distribution whose mean μ and variance σ^2 we wish to infer. In this case,

$$L = \prod_i \left(\frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x_i-\mu)^2}{2\sigma^2} \right) \right)^{C_i}$$

where i is an index over the outcomes of O, C_i is the multiplicity of outcome i in O, and x_i is the outcome itself. Taking the logarithm yields

$$\ell = \sum_i \log\left(\frac{C_i}{\sigma\sqrt{2\pi}} \right) - \frac{C_i(x_i-\mu)^2}{2\sigma^2} \tag{3.1}$$

Now suppose that we expect the mean μ to vary as 2 input parameters d_i and δ_i (which are specified with each observation in O) vary so that

$$\mu = \alpha_0 + d_i\alpha_1 + \delta_i\alpha_2 + \delta_i d_i \alpha_3 \tag{3.2}$$

We can now take the derivative of ℓ with respect to α_0, α_1, α_2, α_3, and σ. Since these are all continuous variables, the derivatives exist and we may solve the polynomial system to obtain the best parameter values. $\diamond$

Several distributions whose probability density functions are compatible with this approach are presented in §3.2.1, §3.2.2, and §3.2.3.

3.1.2 Mixture model and discrete-parameter inference

The second class of problems works when ℓ is composed of a sum of logarithms, as in Equation 1.1. This typically arises when dealing with mixture models and the parameters inferred are the relative frequencies with which observations are drawn from different distributions.

Inferring model parameters other than mixture frequencies with this approach does not involve taking the derivative of the log-likelihood function with respect to distribution parameters so even integer-valued parameters may be inferred. However, the number of variables in the resulting system is related to the number of parameters and the number of discrete values each parameter takes on so it is difficult to consider a large number of parameter values. To be specific, if we are mixing M models and have q parameters to vary, with each parameter taking on n_i distinct values ($i \in \{1, 2, \ldots, q\}$), a total of

$$N_v = (M-1) + \sum_{i=1}^{q} (n_i - 1)$$

variables will be required. Consider the case where we wish to mix a normal and a binomial distribution with 3 possible values for μ and 2 possible values for each of σ, n, and p. Then $N_v = (2-1) + (3-1) + 3(2-1) = 6$.

This pattern of variables is evinced because we consider the choice of a value for each parameter i to be independent of the choice of any another. We use q simplices Δ_{n_i-1} to model the frequency with which each parameter is chosen and one additional simplex Δ_{M-1} to model the frequency with which a particular distribution is chosen. This in turn manifests itself as the logarithm of a product which simplifies to a sum of logarithms.

While multiplicative distributions are not common, it is possible they could occur. For example, if each observation of an experiment was in fact a realization of several independent processes. In this case, the product of several independent distributions is taken to be a single overall model. If parameters of component distributions in the overall model need to be inferred, this may be accomplished using the same procedure as a single distribution in §3.1.1.

Now that we've outlined the two major modeling techniques (a single distribution or a mixture model), we'll provide more details on probability distributions of interest.

- Single "primitive" distribution or product of "primitive" distributions:
  - Distributions of interest: Normal, Log-Normal, Poisson, binomial, Weibull.

  - Adding complexity (and flexibility) by making parameters functions of input variables.

- Shock response spectra
  - Physical phenomena involved.
  - Administrivia: normalizing and scaling inputs and outputs.
  - SRS example. MLE of Gaussian w/ mu a function of 1 or 2 vars.

- Generic mixture model formulation.
  - Terminology, notation, generic formula.
  - SRS example.

3.2 Distributions of interest

One advantage of performing MLE in this setting is that it is possible to infer parameters for distributions other than Gaussians. While some parameters (notably integer-values parameters in discrete probability distributions) cannot be inferred using the method of §3.1.1 because the gradient is not well-defined, this still leaves many useful parameters. The following sections present the total differentials of the log-likelihood for several distributions of interest so that modelers can derive their own MLE equations. In the simplest case, this simply involves setting each term of the total differential to 0 and solving. For models that are products of distributions or whose parameters should be modeled as polynomials in terms of other variables of interest, more work is required.

3.2.1 Poisson distribution

The Poisson distribution is discrete with semi-infinite support and has a probability mass function

$$m_{\text{Poisson}}(k;\lambda) = \frac{e^{-\lambda}\lambda^k}{k!}$$

where $k \in \{1, 2, \ldots\}$ and the differential of its log is

$$\mathrm{d}(\log(m_{\text{Poisson}}(k;\lambda))) = \left(-1 + \frac{k}{\lambda}\right)\mathrm{d}\lambda.$$

Note that k may not be inferred using our MLE approach but this may not generally be necessary because k will typically be used for a model output.

This distribution is useful for modeling the probability of an event occurring when the rate of events λ is constant between observations.

3.2.2 Binomial distribution

The binomial distribution is discrete and compact and has a probability mass function

$$m_{\text{binomial}}(k;p,N) = \binom{N}{k} p^k (1-p)^k$$

where $k \in \{0,1,2,\ldots,N\}$ and the differential of its log is

$$\mathrm{d}(\log(m_{\text{binomial}}(k;p,N))) = \left(\frac{k}{p} - \frac{k}{1-p}\right)\mathrm{d}p.$$

Note that N may not be inferred using our MLE approach but this may not generally be necessary. Because k will typically be used for a model output (whose likelihood we are estimating) that must be normalized by N, inferring N would involve additional work should $k = k(N)$.

This distribution is useful for modeling the probability of an event occurring when the chance of an event between observations is equal to the chance that k choices (with replacement) from a pool of n are identical.

3.2.3 Normal (Gaussian) distribution

The normal distribution is continuous with infinite support and has a probability density function

$$d_{\text{normal}}(x;\mu,\sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$$

where $x \in \mathbb{R}$ and the differential of its log is

$$\mathrm{d}(\log(d_{\text{normal}}(x;\mu,\sigma))) = \left(-\frac{1}{\sigma} - \frac{(x-\mu)^2}{\sigma^3}\right)\mathrm{d}\sigma + \left(\frac{x-\mu}{\sigma^2}\right)\mathrm{d}\mu.$$

You may not infer both x and μ for the obvious reason that the gradient will not produce a system of equations with full rank.

It is useful for modeling the probability of an event occurring when many small, independent, additive effects contribute to the occurrence of an event. If the effects are multiplicative and not additive, then a log-normal distribution should be used.

3.3 Modeling Shock Response Spectra

We will be analyzing 2 types of data: experimental data and simulation data. If we attempt to summarize simulation data with a statistical model, the variance present in the model will be due to the model's inability to completely capture the physics being simulated because the model is

either incomplete (meaning it does not embody the physical model that the simulation does at the level of detail of the simulation) or incorrect (meaning that the statistical model contains an expression which contradicts the physical model in the simulation) or both.

When summarizing experimental data with a statistical model, the variance present in the model may be due to its inability to capture the physics of the experiment (again, this may be because the model is incomplete, incorrect, or both), or because of aleatoric or epistemic uncertainty in the experiment itself.

As a final case to consider, we may attempt to use statistical models to analyze the difference between simulations and experimental data. These statistical models might contain variance due to any of the above plus differences between the physical phenomena driving the simulation and those dominating the experimental results and simulation calibration error.

3.3.1 MLE of SRS

As a first model, consider an ensemble of runs used to verify physics packages by simulating the response of an elastic bar to an impact load. The initial displacement of the bar δ and the distance from the end of the bar where the load was applied d are both varied and the shock response spectra (SRS, first developed by Biot [Bio33]) is measured. The first few peak frequencies and magnitudes of the SRS are determined. Theoretical results as well results from several simulation packages are obtained and yield a distribution of magnitude and frequency values.

We expect the magnitude M of the first SRS peak to decrease linearly with distance d and to increase linearly with the initial displacement δ. Because of truncation error and other parameters which we do not take into account the SRS magnitudes observed are a distribution of values.

Equation 3.1 thus seems like a good candidate for MLE. The resulting set of equations is

$$
\begin{aligned}
0 &= -\left(\sum_i C_i M_i\right) + \left(\sum_i C_i\right)\alpha_0 + \left(\sum_i C_i d_i\right)\alpha_1 + \left(\sum_i C_i \delta_i\right)\alpha_2 + \left(\sum_i C_i d_i \delta_i\right)\alpha_3 \\
0 &= -\left(\sum_i C_i d_i M_i\right) + \left(\sum_i C_i d_i\right)\alpha_0 + \left(\sum_i C_i d_i^2\right)\alpha_1 + \left(\sum_i C_i d_i \delta_i\right)\alpha_2 + \left(\sum_i C_i d_i^2 \delta_i\right)\alpha_3 \\
0 &= -\left(\sum_i C_i \delta_i M_i\right) + \left(\sum_i C_i \delta_i\right)\alpha_0 + \left(\sum_i C_i d_i \delta_i\right)\alpha_1 + \left(\sum_i C_i \delta_i^2\right)\alpha_2 + \left(\sum_i C_i d_i \delta_i^2\right)\alpha_3 \\
0 &= -\left(\sum_i C_i d_i \delta_i M_i\right) + \\
&\quad \left(\sum_i C_i d_i \delta_i\right)\alpha_0 + \left(\sum_i C_i d_i \delta_i\right)\alpha_1 + \left(\sum_i C_i d_i \delta_i^2\right)\alpha_2 + \left(\sum_i C_i d_i^2 \delta_i^2\right)\alpha_3 \\
0 &= \left(\sum_i C_i M_i^2\right) - \left(\sum_i C_i\right)\sigma^2 \\
&\quad -2\left[\left(\sum_i C_i M_i\right)\alpha_0 + \left(\sum_i C_i M_i d_i\right)\alpha_1 + \left(\sum_i C_i M_i \delta_i\right)\alpha_2 + \left(\sum_i C_i M_i d_i \delta_i\right)\alpha_3\right] + \\
&\quad \left(\sum_i C_i\right)\alpha_0^2 + \left(\sum_i C_i d_i^2\right)\alpha_1^2 + \left(\sum_i C_i \delta_i^2\right)\alpha_2^2 + \left(\sum_i C_i d_i^2 \delta_i^2\right)\alpha_3^2 + \\
&\quad 2\left[\left(\sum_i C_i d_i\right)\alpha_0\alpha_1 + \left(\sum_i C_i d_i \delta_i\right)\alpha_1\alpha_2 + \left(\sum_i C_i d_i \delta_i^2\right)\alpha_2\alpha_3 + \right. \\
&\quad \left.\left(\sum_i C_i \delta_i\right)\alpha_0\alpha_2 + \left(\sum_i C_i d_i^2 \delta_i\right)\alpha_1\alpha_3 + \left(\sum_i C_i d_i \delta_i\right)\alpha_0\alpha_3\right]
\end{aligned}
$$

Here, i is an index into a set of all distinct combinations of (M, d, δ) that occur in the ensemble of simulations and C_i is the number of times each distinct combination is observed. The system of equations has 5 variables and 5 unknowns. The first 4 equations have 5 monomials each while the final equation has 16 monomials. The ensemble of runs contains 2660 observations with 910 unique outcomes and this system results

$$
\begin{aligned}
0 &= -2.5422\times10^{7}+2.6600\times10^{3}\alpha_0+7.9800\times10^{3}\alpha_1+3.9900\times10^{1}\alpha_2+1.1970\times10^{2}\alpha_3\\
0 &= -7.5473\times10^{7}+7.9800\times10^{3}\alpha_0+2.5796\times10^{4}\alpha_1+1.1970\times10^{2}\alpha_2+3.8694\times10^{2}\alpha_3\\
0 &= -4.2377\times10^{5}+3.9900\times10^{1}\alpha_0+1.1970\times10^{2}\alpha_1+6.6500\times10^{-1}\alpha_2+1.9950\times10^{0}\alpha_3\\
0 &= -1.2580\times10^{6}+1.1970\times10^{2}\alpha_0+3.8694\times10^{2}\alpha_1+1.9950\times10^{0}\alpha_2+6.4490\times10^{0}\alpha_3\\
0 &= 2.8128\times10^{11}-2.6600\times10^{3}\sigma^2\\
&\quad -5.0844\times10^{7}\alpha_0-1.5095\times10^{8}\alpha_1-8.4753\times10^{5}\alpha_2-2.5160\times10^{6}\alpha_3+\\
&\quad 2.6600\times10^{3}\alpha_0^2+2.5796\times10^{4}\alpha_1^2+6.6500\times10^{-1}\alpha_2^2+6.4490\times10^{0}\alpha_3^2+\\
&\quad 1.5960\times10^{4}\alpha_0\alpha_1+2.3940\times10^{2}\alpha_1\alpha_2+3.9900\times10^{0}\alpha_2\alpha_3+\\
&\quad 7.9800\times10^{1}\alpha_0\alpha_2+7.7388\times10^{2}\alpha_1\alpha_3+2.3940\times10^{2}\alpha_0\alpha_3
\end{aligned}
$$

References

[Bak77] Alan Baker. The theory of linear forms in logarithms. 1977.

[Ber03] Daniel J. Bernstein. Computing logarithm intervals with the arithmetic-geometric mean iterations. *available from* `http://cr.yp.to/papers.html` , 2003.

[Bio33] M. Biot. Theory of elastic systems vibrating under transient impulse with an application to earthquake-proof buildings. *Proceedings of the National Academy of Science*, 19:262–268, 1933.

[BRS08] Frederic Bihan, J. Maurice Rojas, and Casey Stella. Faster real feasibility via circuit discriminants, 2008. submitted for publication.

[DFS07] Alicia Dickenstein, Eva Maria Feichtner, and Bernd Sturmfels. Tropical discriminants. *J. Amer. Math. Soc.*, 20:1111–1133, 2007.

[DRRS07] Alicia Dickenstein, J. Maurice Rojas, Korben Rusek, and Justin Shih. Extremal real algebraic geometry and $\mathscr{A}$-discriminants. *Moscow Mathematical Journal*, 7(3):425–452, 2007.

[DS08] Bernd; Drton, Mathias; Sturmfels and Seth Sullivant. *Lectures on Algebraic Statistics*. Birkhäuser, Basel, 2008.

[GKZ94] Israel Moseyevitch Gel'fand, Misha M. Kapranov, and Andrei V. Zelevinsky. *Discriminants, Resultants and Multidimensional Determinants*. Birkhäuser, Boston, 1994.

[Kap91] Misha Kapranov. A characterization of $\mathscr{A}$-discriminantal hypersurfaces in terms of the logarithmic gauss map. *Mathematische Annalen*, 290:277–285, 1991.

[MR04] Gregorio Malajovich and J. Maurice Rojas. High probability analysis of the condition number of sparse polynomial systems. *Theoretical Computer Science, special issue on algebraic and numerical algorithms*, 315(2–3):525–555, May 2004.

[Nes03] Yuri Nesterenko. Linear forms in logarithms of rational numbers. *Lecture Notes in Math.*, 1819:53–106, 2003.

[PS05] Lior Pachter and Bernd Sturmfels. *Algebraic Statistics for Computational Biology*. Cambridge University Press, 2005.

[PT05] Mikael Passare and August Tsikh. Amoebas: their spines and their contours. *Idempotent mathematics and mathematical physics, Contemp. Math.*, 377:275–288, 2005.

[PW01] Eva; Pistone, Giovanni; Riccomagno and Henry Wynn. *Algebraic Statistics: Computational Commutative Algebra in Statistics*. Chapman & Hall/CRC, 2001.

[Sma00] Steve Smale. Mathematical problems for the next century. *Mathematics: Frontiers and Perspectives*, pages 271–294, 2000.

[SW05] Andrew J. Sommese and Charles W. Wampler. *The Numerical Solution of Systems of Polynomials Arising in Engineering and Science*. World Scientific Publishing, 2005.

DISTRIBUTION:

1	J. Maurice Rojas Department of Mathematics Texas A&M University 3368 TAMU College Station, TX 77843	
2	MS 9159	Philippe P. Pébay, 8963
2	MS 9159	David Thompson, 8963
2	MS 9018	Central Technical Files, 8944
1	MS 0899	Technical Library, 9536

Sandia National Laboratories