8

Introduction to numerical algebraic geometry

Andrew J. Sommese[1⋆], Jan Verschelde[2⋆⋆], and Charles W. Wampler[3]

[1] Department of Mathematics, University of Notre Dame, Notre Dame, IN 46556-4618, USA, sommese@nd.edu, http://www.nd.edu/~sommese

[2] Department of Mathematics, Statistics, and Computer Science, University of Illinois at Chicago, 851 South Morgan (M/C 249), Chicago, IL 60607-7045, USA, jan@math.uic.edu, jan.verschelde@na-net.ornl.gov, http://www.math.uic.edu/~jan

[3] General Motors Research and Development, Mail Code 480-106-359, 30500 Mound Road, Warren, MI 48090-9055, USA, Charles.W.Wampler@gm.com

Summary. In a 1996 paper, Andrew Sommese and Charles Wampler began developing a new area, "Numerical Algebraic Geometry", which would bear the same relation to "Algebraic Geometry" that "Numerical Linear Algebra" bears to "Linear Algebra".

To approximate all isolated solutions of polynomial systems, numerical path following techniques have been proven reliable and efficient during the past two decades. In the nineties, homotopy methods were developed to exploit special structures of the polynomial system, in particular its sparsity. For sparse systems, the roots are counted by the mixed volume of the Newton polytopes and computed by means of polyhedral homotopies.

In Numerical Algebraic Geometry we apply and integrate homotopy continuation methods to describe solution components of polynomial systems. In particular, our algorithms extend beyond just finding isolated solutions to also find all positive dimensional solution sets of polynomial systems and to decompose these into irreducible components. These methods can be considered as symbolic-numeric, or perhaps rather as numeric-symbolic, since numerical methods are applied to find integer results, such as the dimension and degree of solution components, and via interpolation, to produce symbolic results in the form of equations describing the irreducible components.

Applications from mechanical engineering motivated the development of Numerical Algebraic Geometry. The performance of our software on several test problems illustrates the effectiveness of the new methods.

⋆ This material is based upon work supported by the National Science Foundation under Grant No. 0105653; and the Duncan Chair of the University of Notre Dame.

⋆⋆ This material is based upon work supported by the National Science Foundation under Grant No. 0105739 and Grant No. 0134611.

8.0 Introduction

The goal of this chapter is to provide an overview of the main ideas developed so far in our research program to implement numerical algebraic geometry, initiated in [SW96].

We are concerned with numerically solving polynomial systems. While the homotopy continuation methods of the past were limited to approximating only the isolated roots, we developed tools to describe all positive dimensional irreducible components of the solution set of a polynomial system. In particular, our algorithms produce for every irreducible component a *witness set*, whose cardinality equals the degree of the component, as this set is obtained by intersecting the component with a general linear space of complementary dimension. A point of a witness set corresponds to what is known in algebraic geometry as a generic point. Our main results [SV00, SVW01a, SVW01b, SVW01c, SVW02c, SVW02b, SVW02a, SVW03, SVW, SVW04] can be summarized in four items:

1. In [SV00] we presented a cascade of homotopies (extended in [SVW]) to find candidate witness points for every component of the solution set. Separating the junk from the candidate witness points was done in [SVW01a], where factorization methods based on interpolation implemented a numerical irreducible decomposition. The use of central projections and a homotopy membership test to filter junk were the improvements of [SVW01b].
2. The treatment of high-degree components and components of multiplicity greater than one can present numerical challenges. The use of monodromy [SVW01c] followed by the validation by the linear trace [SVW02c] enabled us to deal with high degree components of multiplicity one, using only machine floating point numbers. In [SVW02b], we presented an approach to tracking paths on sets of multiplicity greater than one, which in theory makes the algorithm for irreducible decomposition completely general, although in practice this portion of the framework needs further refinement. However, for the case of the factorization of a single multivariate polynomial, we can use differentiation to reduce the treatment of higher multiplicity components to nonsingular path tracking, as we described in [SVW04]. This addresses an open problem in symbolic-numeric computing: the factorization of multivariate polynomials with approximate coefficients [Kal00].
3. Our new homotopy algorithms have been implemented and tested using the path trackers in the software package PHCpack [Ver99a]. In [SVW03] we outlined the new tools in PHCpack and described a simple interface to Maple. Our software found the degrees of all irreducible components of the cyclic 8 and 9 roots problems, which previously could only be done via Gröbner bases (and only by the very best implementation [Fau99]).
4. Polynomial systems with positive dimensional components occur naturally when designing mechanical devices which permit motion. We inves-

tigated a special case of a moving platform, discovering through a numerical irreducible decomposition [SVW02c] a component not reported by experts [HK00]. This and other applications of our tools to systems coming from mechanical design are described in [SVW02a].

In this chapter we will introduce these results, first explaining homotopy methods for isolated solutions. We can only mention some recent and exciting new developments in fields related to numerical algebraic geometry: numerical Schubert calculus ([HSS98], [HV00], [LWW02], [SS01], [VW02]) and numerical jet geometry [RSV02].

8.1 Homotopy continuation methods – an overview

Homotopy continuation methods operate in two stages. Firstly, homotopy methods exploit the structure of the system $f(\mathbf{x}) = 0$ to find a root count and to construct a start system $g(\mathbf{x}) = 0$ that has exactly as many regular solutions as the root count. This start system is embedded in the *homotopy*

$$h(\mathbf{x}, t) = \gamma(1-t)g(\mathbf{x}) + tf(\mathbf{x}) = 0, \quad t \in [0,1], \tag{8.1}$$

with $\gamma \in \mathbb{C}$ a random number. Secondly, as t moves from 0 to 1, numerical continuation methods trace the paths that originate at the solutions of the start system towards the solutions of the target system. The good properties we expect from a homotopy are (borrowed from [Li97, Li03]):

1. (*triviality*) The solutions for $t = 0$ are trivial to find.

2. (*smoothness*) No singularities along the solution paths occur (because of γ).

3. (*accessibility*) An isolated solution of multiplicity m is reached by exactly m paths.

Continuation or path-following methods are standard numerical techniques ([AG90a, AG93, AG97], [Mor87], [Wat86, Wat89]) to trace the solution paths defined by the homotopy using *predictor-corrector* methods. The smoothness property of complex polynomial homotopies implies that paths never turn back, so that during correction the parameter t stays fixed, which simplifies the set up of path trackers. The adaptive step size control determines the step length while enforcing quadratic convergence in Newton's method to avoid path crossing (see also [KX94] for the application of interval methods to control the step size). At the end of the path, end games ([HV98], [MSW91, MSW92a, MSW92b], [SWS96]) deal with diverging paths and paths leading to singular roots.

Following [HSS98], we say that a homotopy is *optimal* if every path leads to one solution. The classification in Table 8.1 (from [Ver99b]) contains key words for three classes of polynomial systems for which optimal homotopies are available in PHCpack [Ver99a]. These homotopies have no diverging paths for generic instances of polynomial systems in their class.

system	model	theory	space	
dense	highest degrees	Bézout	$\mathbb{P}^n$	projective
sparse	Newton polytopes	Bernshteĭn	$(\mathbb{C}^*)^n$	toric
determinantal	localization posets	Schubert	G_{mr}	Grassmannian

Table 8.1. Key words of the three classes of polynomial systems.

The earliest applications of homotopies for solving polynomial systems ([CMPY79], [Dre77], [GZ79], [GL80], [Li83], [LS87] [Mor83], [Wri85], [Zul88]) belong to the dense class, where the number of paths equals the product of the degrees in the system. Multi-homogeneous homotopies were introduced in [MS87b, MS87a] and applied in [WMS90, WMS92], see also [Wam92]. Similar are the random product homotopies [LSY87a, LSY87b], see also [Li87] and [LW91]. Methods to construct linear-product start systems were introduced in [VH93], and extended in [VC93, VC94], [LWW96], and [WSW00]. A general approach to exploit product structures was developed in [MSW95].

Almost all systems have fewer terms than allowed by their degrees. Implementing constructive proofs of Bernshteĭn's theorems [Ber75], polyhedral homotopies were introduced in [HS95] and [VVC94] to solve sparse systems more efficiently. These methods provided ways to start cheater's homotopies ([LSY89], [LW92]) and special instances of coefficient-parameter polynomial continuation ([MS89, MS90]). The root count requires the calculation of the mixed volume[4], for which a lift-and-prune approach was presented in [EC95]. Exploitation of symmetry was studied in [VG95] and the dynamic lifting of [VGC96] led to incremental polyhedral continuation. See [Ver00] for a Toric Newton. Extensions to count all affine roots (also those with zero components) were proposed in [EV99], [GLW99], [HS97b], [LW96], [Roj94, Roj99b], and [RW96]. Very efficient calculations of mixed volumes are described in [DKK03], [GL00, GL03], [KK03b], [LL01], and [TKF02].

Determinantal systems (with equations like $\det(A|X) = 0$) arise in problems of enumerative geometry. The homotopies in numerical Schubert calculus first appeared explicitly in [HSS98], originating from questions in real enumerative geometry [Sot97a, Sot97b]. While real enumerative geometry [Sot03] is interesting on its own, these homotopies solve the pole placement problem ([Byr89], [RRW96, RRW98], [Ros94], [RW99]) in control theory. Recent improvements and applications can be found in [HV00], [LWW02], [SS01], and [VW02].

We end this section noting that homotopies have a wider application range than "just" solving polynomial systems, see for instance [Wat02] for a survey, [WBM87], and [WSM$^+$97] for a description of HOMPACK. The

[4] The mixed volume was nicknamed in [CR91] as the BKK bound to honor Bernshteĭn [Ber75], Kushnirenko [Kus76], and Khovanskiĭ [Kho78b].

speedup of continuation methods on multi-processor machines has been addressed in [ACW89, CARW93, HW89].

8.2 Homotopies to approximate all isolated solutions

We first prove the regularity and boundedness of the solution paths defined by homotopies, before surveying path following techniques. We obtain more efficient homotopies by exploiting product structures and using Newton polytopes to model the sparsity of the system.

8.2.1 Regularity and boundedness of solution paths

To illustrate how homotopy methods work, let us consider a simple example of solving two quadrics:

$$f(x,y) = \begin{pmatrix} x^2 + 4y^2 - 4 \\ 2y^2 - x \end{pmatrix}.$$

To solve $f(x,y) = 0$, we match it with a start system of two easily solved quadrics:

$$g(x,y) = \begin{pmatrix} x^2 - 1 \\ y^2 - 1 \end{pmatrix},$$

with which we form the following homotopy:

$$h(x,y,t) = \begin{pmatrix} x^2 - 1 \\ y^2 - 1 \end{pmatrix}(1-t) + \begin{pmatrix} x^2 + 4y^2 - 4 \\ 2y^2 - x \end{pmatrix} t. \tag{8.2}$$

At $t = 1$, $h(x,y,t=1) = 0$ is $f(x,y) = 0$, the system we wish to solve while at $t = 0$, $h(x,y,t=0) = 0$ is the start system $g(x,y) = 0$ we can easily solve. As we usually move t from 0 to 1 when we solve the system, we may view the movement of t from 1 to 0 as a degeneration of the system, i.e., we deform the general hypersurfaces into degenerate products of hyperplanes.

But does this work? We will see in a moment that it does not, but that there is a simple maneuver that fixes the trouble once and for all. For numerical solving, we would need the solution paths to be free of singularities. A singularity occurs where the Jacobian matrix J_h of the homotopy $h(x,y,t) = 0$ has a zero determinant. The singularities along the solution paths are solutions of the system

$$\begin{cases} h(x,y,t) = \mathbf{0} \\ \det(J_h(x,y,t)) = 0 \end{cases} \quad \text{where} \quad J_h = \begin{bmatrix} 2x & 8yt \\ -t & 2y + 2yt \end{bmatrix}. \tag{8.3}$$

If this "discriminant system" has any roots with $t \in [0,1)$, there is at least one homotopy solution path with singularities. To explore this situation, let's solve this system by elimination. This is not a step that we normally perform

in the course of solving $f(x) = 0$, but we do it here to reveal the flaw in the naive homotopy of (8.2) and to illustrate how we fix the flaw. To solve this discriminant system, we will eliminate from the system the variables x and y to obtain one polynomial in the continuation parameter t. The roots of this polynomial define the singularities along the solution paths.

While there are many ways to perform this elimination, we let Maple compute a lexicographical Gröbner basis of the discriminant system. Below are the Maple commands, to save space we suppressed most of the output.

```
> f := [x^2 + 4*y^2 - 4,2*y^2- x];      # target system
> g := [x^2 - 1, y^2 - 1];              # start system
> h := t*f + (1-t)*g;                   # the homotopy
> eh := expand(h);                      # expanded homotopy
> jh := matrix(2,2,                     # Jacobian matrix
         [[diff(eh[1],x),diff(eh[1],y)],
          [diff(eh[2],x),diff(eh[2],y)]]);
> sys := [eh[1],eh[2],      # discriminant system solved by
          linalg[det](jh)]; # pure lex Groebner basis in gb
> gb := grobner[gbasis](sys,[x,y,t],plex);
> gb[nops(gb)];             # discriminant polynomial
                    3       5       4      2       7       6
      -1 + t + 10 t  + 29 t  + 13 t  - 5 t  + 12 t  + 21 t
```

As the degree of this "discriminant polynomial" is seven, we have seven roots:

```
> fsolve(gb[nops(gb)],t,complex);       # numerical solving
  -.8818537646 - .9177002576 I, -.8818537646 + .9177002576 I,
  -.2011599690 - .8877289373 I, -.2011599690 + .8877289373 I,
.006853764567 - .3927967328 I, .006853764567 + .3927967328 I,
   .4023199381
```

We are troubled by the root around 0.4, because, as t moves from 0 to 1, we will encounter a singularity. So our homotopy in (8.2) does not work!

We can fix this problem by the choice of a random constant $\gamma = e^{\theta\sqrt{-1}}$, for some random angle θ. Now, consider the homotopy

$$h(x, y, t) = \gamma \begin{pmatrix} x^2 - 1 \\ y^2 - 1 \end{pmatrix} (1 - t) + \begin{pmatrix} x^2 + 4y^2 - 4 \\ 2y^2 - x \end{pmatrix} t. \tag{8.4}$$

The random choice of γ will cause all roots of the discriminant polynomial to lie outside the interval $[0, 1)$. That $t = 0$ is excluded is obvious (because the start system has only regular roots), but at $t = 1$ we may find singular solutions of the given system f.

Exercise 8.2.1. Modify the homotopy in the sequence of Maple commands above taking `h := t*f + (1+I)*(1-t)*g;` and verify that none of the roots

of the discriminant polynomial is real. The choice of γ as $1 + \sqrt{-1}$ does not give the Gröbner package of Maple a hard time. If Maple is unavailable, then another computer algebra system should do just as well.

The above example illustrates the general idea behind the regularity of solution paths defined by a homotopy. The main theorem of elimination theory says that the projection of an algebraic set in complex projective space is again an algebraic set. Consider the discriminant system as a polynomial system in $\mathbf{x}, t$, and γ. If we eliminate $\mathbf{x}$, we obtain a polynomial in t and γ. This polynomial does not vanish entirely as the start system (at $t = 0$) has no singular roots. Thus it has only finitely many roots for general γ. Furthermore, a random complex choice of γ will insure that all those roots miss the interval $[0, 1)$. A schematic (as in [Mor87]) illustrating what cannot and what can happen is in Figure 8.1.

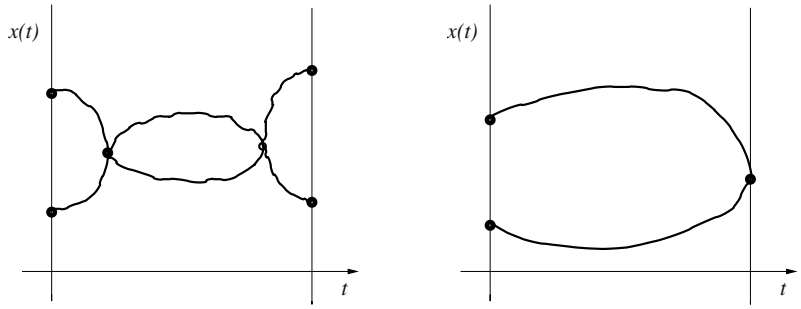

Fig. 8.1. By a random choice of a complex constant γ, singularities will not occur for all $t \in [0, 1)$ as on the left, but they may occur at the end, for $t = 1$.

The same random constant γ ensures that all paths stay bounded for all $t \in [0, 1)$. By this we mean that no path diverges to infinity for some $t \in [0, 1)$. Equivalently, for all $t \in [0, 1)$, the system $h(\mathbf{x}, t) = \mathbf{0}$ has no solutions at infinity (see Figure 8.2). To see this, invoke a homogeneous coordinatetransformation introducing one extra coordinate, and consider the system in projective space. That is, consider the homogenized system $H(X, Y, Z, t) = 0$ obtained by clearing Z from denominators in the expression $h(X/Z, Y/Z, t) = 0$. Now, instead of the discriminant system of (8.3) our concern is the system

$$\begin{cases} H(X, Y, Z, t) = 0 \\ Z = 0 \end{cases}$$

Since h is homogeneous in X, Y, Z, the solutions live in projective space, which we can restate to say that all solutions to $H(X, Y, 0, t) = 0$ must either satisfy $H(X/Y, 1, 0, t) = 0$ or $H(1, Y/X, 0, t) = 0$ (or both, if neither X or Y is zero). Either of these is a system of two polynomials in two variables and γ and so

one can again apply elimination and see that, except for special choices of γ, there will be no solutions at infinity for $t \in [0, 1)$.

Note that if the polynomials in the start system $g(x, y) = 0$ have lower degrees than their counterparts in $f(x, y) = 0$, then $H(X, Y, Z, t) = 0$ could have solutions at infinity for $t = 0$. By matching the degrees of the polynomials in g and f, we avoid this, which is key in proving the third property of a good homotopy: accessibility.

Exercise 8.2.2. Consider the homotopy

$$h(x, y, t) = \left(\begin{cases} x^2 - 1 = 0 \\ y^2 - 1 = 0 \end{cases} \right) (1 - t) + \left(\begin{cases} y^2 - 1 = 0 \\ x^2 - 3 = 0 \end{cases} \right) t.$$

For which values of t do we have diverging paths? Show that with a random complex constant γ in $h(x, y, t) = \mathbf{0}$ (as in (8.4)) there are no divergent paths.

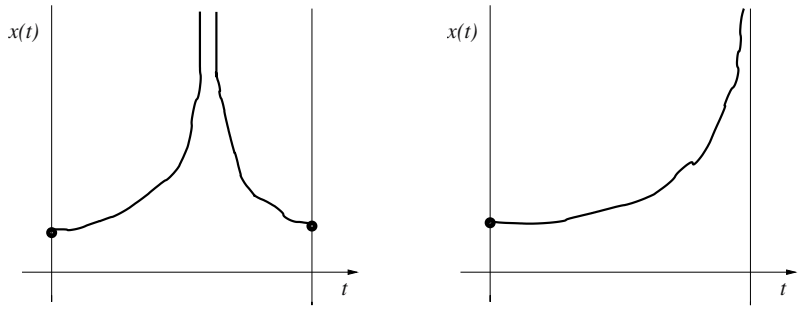

Fig. 8.2. By a random choice of a complex constant γ, divergence will not occur for all $t \in [0, 1)$ as on the left, but may occur at the end, for $t = 1$.

To understand why the homotopy has the accessibility property (defined in Section 8.1), consider that whenever the number of equations is equal to the number of variables $\mathbf{x}$, continuity implies that an isolated root at $t = 1$ must be approached by at least one isolated root as $t \to 1$. Since there are no singularities or solutions at infinity for t in $[0, 1)$, we can carry this argument backwards all the way to $t = 0$, where we know we are starting with all the solutions of the homotopy.

The arguments described above can be found in [BCSS98], see also [LS87].

8.2.2 Path following techniques

Consider any homotopy $h_k(x(t), y(t), t) = 0$, $k = 1, 2$. Since we are interested to see how x and y change as t changes, we apply the operator $\frac{\partial}{\partial t}$ on the homotopy. Via the chain rule, we obtain

$$\frac{\partial h_k}{\partial x}\frac{\partial x}{\partial t} + \frac{\partial h_k}{\partial y}\frac{\partial y}{\partial t} + \frac{\partial h_k}{\partial t} = 0, \quad k = 1, 2.$$

Denote $\Delta x := \frac{\partial x}{\partial t}$ and $\Delta y := \frac{\partial y}{\partial t}$. For fixed t (after incrementing $t := t + \Delta t$), for $k = 1, 2$, we solve the linear system

$$\begin{bmatrix} \frac{\partial h_1}{\partial x} & \frac{\partial h_1}{\partial y} \\ \frac{\partial h_2}{\partial x} & \frac{\partial h_2}{\partial y} \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = - \begin{bmatrix} \frac{\partial h_1}{\partial t} \\ \frac{\partial h_2}{\partial t} \end{bmatrix}$$

and obtain $(\Delta x, \Delta y)$, the tangent to the path. For some step size $\lambda > 0$, the updates $x := x + \lambda \Delta x$ and $y := y + \lambda \Delta y$ give the Euler predictor.

To avoid solving a linear system at each predictor step, we may use a secant predictor. A secant predictor is less accurate and will require more corrector steps, but the total amount of work for the prediction can be less. Cubic interpolation, using the tangent vectors at two points along the path, leads to the Hermite predictor. See Figure 8.3 for a comparison.

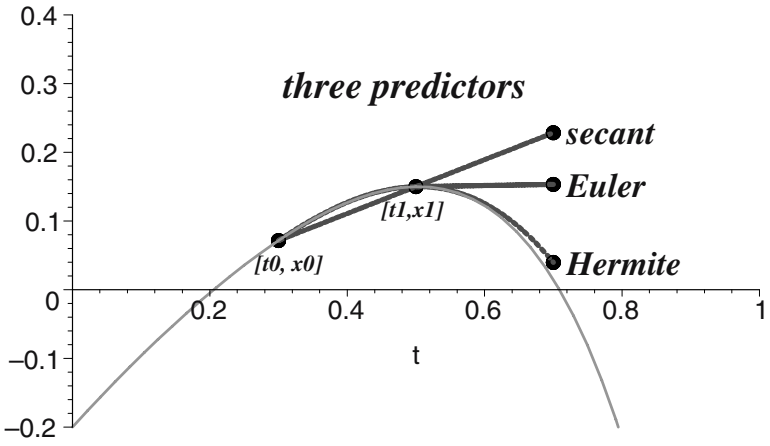

Fig. 8.3. Three predictors: secant, Euler, and Hermite.

The *predictor* delivers at each step of the method a new value of the continuation parameter and predicts an approximate solution of the corresponding new system in the homotopy. Then, the predicted approximate solution is corrected by applying the *corrector*, e.g., by Newton's method. With a good homotopy, the solution paths never turn back as t increases. Therefore, the continuation parameter can remain fixed while correcting the predicted solution. This leads to so-called *increment-and-fix* path following methods. In practice, determining the step length during the prediction stage is done by a hit-or-miss method, which can be implemented by means of an *adaptive step size control*, as done in the algorithm below.

Algorithm 8.2.3 Following one solution path by an increment-and-fix predictor-corrector method with an adaptive step size control strategy.

```
Input:  h(x,t), x* ∈ C^n: h(x*,0) = 0,          homotopy and root
        ε > 0, max_it, max_steps,               defines stop criteria
        min_step_size, max_step_size.           for step size control
Output: x*, success if ||h(x*,1)|| ≤ ε.         approximate root at end

t := 0;  k := 0;                                initialization
λ := max_step_size;                             step length
old_t := t;  old_x* := x*                       back up for t and x*
previous_x* := x*;                              previous solution
stop := false;                                  combines stop criteria
while t < 1 and not stop loop
   t := min(1, t + λ);                          secant predictor for t
   x* := x* + λ(x* - previous_x*);              secant predictor for x*
   Newton(h(x,t), x*, ε, max_it, success);      correct with Newton
   if success                                   step size control
    then λ := min(Expand(λ), max_step_size);    enlarge step length
         previous_x* := old_x*;                 go further along path
         old_t := t;  old_x* := x*;             new back up values
    else λ := Shrink(λ);                        reduce step length
         t := old_t;  x* := old_x*;             step back and try again
   end if;
   k := k + 1;                                  augment counter
   stop := (λ < min_step_size)                  1st stop criterion
        or (k > max_steps);                     2nd stop criterion
end loop;
success := (||h(x*,1)|| ≤ ε).                   report success or failure
```

The path following algorithm contains three key ingredients in its loop: the predictor, the corrector and the step size control. The step size λ is controlled by the functions *Shrink* and *Expand* which respectively reduce and enlarge λ, depending on the outcome of the corrector.

The algorithm is still abstract because we did not specify particular values for the constants, such as tolerances on the solutions, minimal and maximal step size, maximum number of iterations of Newton's method, etc.

8.2.3 Homotopies exploiting product structures

A typical homotopy looks as follows:

$$h(\mathbf{x}, t) = \gamma g(\mathbf{x})(1-t) + f(\mathbf{x})t = 0, \quad \gamma \in \mathbb{C},$$

where a random γ ensures the regularity and boundedness of the paths.

In general, for a system $f = (f_1, f_2, \ldots, f_n)$, with $d_i = \deg(f_i)$, we set up a start system $g(\mathbf{x}) = \mathbf{0}$ as follows:

$$g(\mathbf{x}) = \begin{cases} \alpha_1 x_1^{d_1} - \beta_1 = 0 \\ \alpha_2 x_2^{d_2} - \beta_2 = 0 \\ \qquad \vdots \\ \alpha_n x_n^{d_n} - \beta_n = 0 \end{cases}$$

where the coefficients α_i and β_i, for $i = 1, 2, \ldots, n$, are chosen at random in $\mathbb{C}$. Therefore $g(\mathbf{x}) = \mathbf{0}$ has exactly as many regular solutions as the total degree $D = \prod_{i=1}^{n} d_i$. So this homotopy defines D solution paths. The theorem of Bézout (which can be proven constructively via a homotopy) indeed predicts D as the number of solutions *in complex projective space*.

Exercise 8.2.4. Consider the following polynomial system:

$$\begin{cases} x^{108} + 1.1y^{54} - 1.1y = 0 \\ y^{108} + 1.1x^{54} - 1.1x = 0 \end{cases}.$$

This system was constructed by Bertrand Haas [Haa02] who provided with this system a counterexample to the conjecture of Kushnirenko on the number of real roots of sparse systems. Use `phc` (available via [Ver99a]) to determine[5] how many solutions of this system are complex. How many are real?

In almost all applications, the systems have far fewer solutions than the total degree (most solutions lie at infinity and are of no interest). Consider the eigenvalue problem $A\mathbf{x} = \lambda\mathbf{x}$, $A \in \mathbb{C}^{n\times n}$. To make the system square, we can add one general hyperplane to obtain a unique $\mathbf{x}$ for every λ. If we apply Bézout's theorem in a straightforward manner, we consider $A\mathbf{x} = \lambda\mathbf{x}$ as a system of n quadrics and obtain a homotopy with $D = 2^n$ to trace, whereas we know there can be at most n solutions! This is a highly wasteful computation, as $2^n - n$ of our solution paths are certain to diverge to infinity.

Let us examine the smallest nontrivial case: $n = 2$. We consider a general 2-by-2 matrix A and scale the components of the eigenvector with a random hyperplane $c_0 + c_1x_1 + c_2x_2 = 0$. So we look at the system

$$f(x_1, x_2, \lambda) = \begin{cases} a_{11}x_1 + a_{12}x_2 - \lambda x_1 = 0 \\ a_{21}x_1 + a_{22}x_2 - \lambda x_2 = 0 \\ \qquad c_0 + c_1x_1 + c_2x_2 = 0 \end{cases}.$$

To compute the solutions at infinity, we go to homogeneous coordinates, replacing x_1 by x_1/x_0, x_2 by x_2/x_0, and λ by λ/x_0. Clearing denominators:

[5] This may take some time (especially on slower machines)...

$$f(x_0,x_1,x_2,\lambda) = \begin{cases} a_{11}x_0x_1 + a_{12}x_0x_2 - \lambda x_1 = 0 \\ a_{21}x_0x_1 + a_{22}x_0x_2 - \lambda x_2 = 0 \\ c_0x_0 + c_1x_1 + c_2x_2 = 0 \end{cases}.$$

Solutions at infinity are solutions of the homogeneous system with $x_0 = 0$ and not all components equal to zero. If $\lambda = 0$, then $(x_0,x_1,x_2,\lambda) = (0,1,-c_1/c_2,0)$ represents one point at infinity. If $\lambda \neq 0$, then the other solution at infinity is represented by $(x_0,x_1,x_2,\lambda) = (0,0,0,1)$. So we found where two of the four paths are diverging to.

Now we embed our problem in *multi-projective* space: $\mathbb{P} \times \mathbb{P}^2$, separating λ from $\mathbf{x}$. To go to 2-homogeneous coordinates, we replace x_2 by x_2/x_0, x_1 by x_1/x_0 (as before), and λ by λ_1/λ_0 (this is new), clearing denominators:

$$f(x_0,x_1,x_2,\lambda_0,\lambda_1) = \begin{cases} a_{11}\lambda_0x_1 + a_{12}\lambda_0x_2 - \lambda_1x_1 = 0 \\ a_{21}\lambda_0x_1 + a_{22}\lambda_0x_2 - \lambda_1x_2 = 0 \\ c_0x_0 + c_1x_1 + c_2x_2 = 0 \end{cases}. \tag{8.5}$$

Looking for roots at infinity of (8.5) we see that $\lambda_0 = 0$ implies $x_1 = 0$, $x_2 = 0$, and thus $x_0 = 0$, so we have no proper solution at infinity with $\lambda_0 = 0$. For the solutions at infinity of (8.5) with $x_0 = 0$, considering (8.5) back in affine coordinates for λ (as λ_0 cannot be zero), we are looking at a homogeneous system of three equations in three unknowns: x_1, x_2, and λ. For general matrices, the trivial zero solution is the only solution. Thus in $\mathbb{P} \times \mathbb{P}^2$, the general eigenvalue problem has no solutions at infinity.

To arrive at a version of Bézout's theorem for polynomial systems over multi-projective spaces, we need to define our root count. Continuing our running example, we record the degrees in λ and $\{x_1,x_2\}$ of every equation in a table. Corresponding to this degree table is a linear-product start system, written in (8.6) in table format.

	$\{\lambda\}$	$\{x_1,x_2\}$
(1)	1	1
(2)	1	1
(3)	0	1

degree table

$\Longleftrightarrow$

	$\{\lambda\}$	$\{x_1,x_2\}$
(1)	$\alpha_{10} + \alpha_{11}\lambda$	$\beta_{10} + \beta_{11}x_1 + \beta_{12}x_2$
(2)	$\alpha_{20} + \alpha_{21}\lambda$	$\beta_{20} + \beta_{21}x_1 + \beta_{22}x_2$
(3)	1	$\beta_{30} + \beta_{31}x_1 + \beta_{32}x_2$

linear-product start system

(8.6)

The coefficients α_{ij} and β_{ij} in (8.6) are randomly chosen complex numbers. Except for a special choice of these numbers, the linear-product start system will always have two regular solutions. We derive a formal root count following the moves we make to solve the linear-product start system:

$$\begin{array}{rccccccccccccccc} B = & 1 & \times & 1 & \times & 1 & + & 1 & \times & 1 & \times & 1 & + & 0 & \times & 1 & \times & 1. \\ & (1)_\lambda & & (2)_{\mathbf{x}} & & (3)_{\mathbf{x}} & & (2)_\lambda & & (1)_{\mathbf{x}} & & (3)_{\mathbf{x}} & & (3)_\lambda & & (1)_{\mathbf{x}} & & (2)_{\mathbf{x}} \end{array} \tag{8.7}$$

The labels in (8.7) show the navigation through the table at the right of (8.6).

Exercise 8.2.5. The matrix polynomial

$$p(\lambda) = A_d\lambda^d + A_{d-1}\lambda^{d-1} + \cdots + A_1\lambda + A_0, \quad A_i \in \mathbb{C}^{n\times n},$$

defines the generalized eigenvalue problem $p(\lambda)\mathbf{x} = \mathbf{0}$. How many generalized eigenvalue-eigenvector pairs can we expect for randomly chosen matrices A_i?

To show that B is an upper bound for the number of isolated solutions of a polynomial system, we show the regularity and boundedness of the solution paths in a typical homotopy, using a linear-product start system.

For many applications (like the eigenvalue problem) it is obvious how best to separate the variables into a partition. But for black-box solvers and systems with no apparent product structure, we need to find that partition which leads to the smallest Bézout number. One strategy is to enumerate all partitions and retain the partition with the smallest Bézout number. While the number of partitions grows faster than 2^n, finding the smallest Bézout number for $n = 8$ by enumeration takes less than a second of CPU time.

Instead of using one partition of the variables to model the product structure of the system, we may use different partitions for different equations, and extend this even further to construct in this way general linear-product start systems. The solving of the start system now involves more work, but we may expect the homotopy to be more efficient. Schematically, a hierarchy of homotopies (and root counting methods) is given in Figure 8.4.

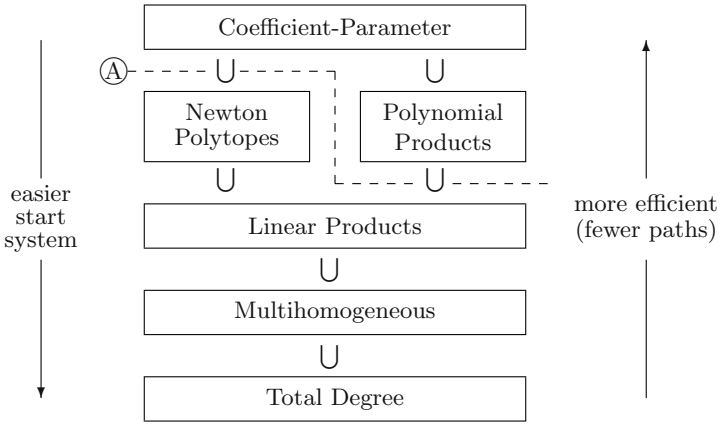

Fig. 8.4. A hierarchy of homotopies. All homotopies below the dashed line A can be done automatically. Above the line, apply special ad-hoc methods or bootstrapping. Homotopies at the bottom of the hierarchy are often used to find solutions for generic instances of parameters in a coefficient-parameter homotopy.

We will not address the "polynomial products" of Figure 8.4 here; for this, see [MSW95]. We introduce the Newton polytopes in the following two sections.

For the relation between Newton polytopes and resultants, see Chapter 7.

8.2.4 Polyhedral homotopies to glue real solutions

The purpose of this section is to introduce Newton polytopes and polyhedral homotopies, but without mixed volumes. So we restrict ourselves to polynomials in one variable. Instead of "just" solving a polynomial in one variable, we consider a different problem:

Input: k distinct monomials in one variable x:
$x^{a_1}, x^{a_2}, \ldots, x^{a_k}$, with $a_i \neq a_j$ for $i \neq j$.
Output: coefficients $c_{a_1}, c_{a_2}, \ldots, c_{a_k}$ such that
$f(x) = c_{a_1}x^{a_1} + c_{a_2}x^{a_2} + \cdots + c_{a_k}x^{a_k}$
has $k-1$ positive real roots.

For example, take $1, x^5, x^7, x^{11}$ as monomials on input. Then the problem is to find c_0, c_5, c_7, and c_{11} such that $f(x) = c_0 1 + c_5 x^5 + c_7 x^7 + c_{11} x^{11}$ has three positive real solutions. We will show that we can reduce this four dimensional problem in that of one dimension, considering the homotopy

$$h(x,t) = t - x^5 + x^7 - x^{11}t = 0, \quad \text{for } t \geq 0.$$

The alternation of signs in the coefficients is a deliberate choice to maximize the number of positive real roots. The Newton polytope of a polynomial is the convex hull of the exponent vectors of those monomials appearing with a nonzero coefficient. The choice of powers of t with each monomial is such that the lower hull of the Newton polytope of h contains among its vertices all exponents of the given monomials, see Figure 8.5.

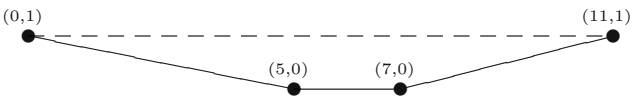

Fig. 8.5. The Newton polytope of the homotopy $h(x,t) = 0$ is spanned by by the exponent vectors of the monomials in h. The lower hull of the Newton polytopes is drawn in solid lines.

At $t = 0$, the homotopy $h(x,0) = -x^5 + x^7 = x^5(-1 + x^2) = 0$ has one positive real root: $x = 1$. The idea is to choose $t = \Delta t > 0$ such that Newton's method applied to $h(x, \Delta t) = 0$ converges quadratically to a positive real root starting at $x = 1$. (Notice that by the fortunate choice of the powers of t in the example, Δt can be chosen arbitrarily large as $h(1,t) \equiv 0$, for any value of t.)

Observe that the monomials in $h(x,0)$ correspond to the lowest middle edge on the lower hull of the Newton polytope of h in Figure 8.5. For every edge of the lower hull of the Newton polytope we will use one homotopy to find one positive real root. Each time, the start system in the homotopy has its

two monomials as vertices of an edge of the lower hull. To find the homotopies with the other two edges, we need to consider the vectors orthogonal to the edges (we call those vectors *inner normals*), see Figure 8.6.

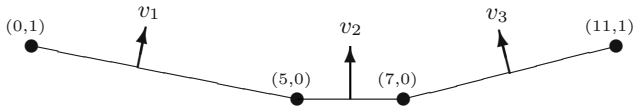

Fig. 8.6. Inner normals $v_1 = (\frac{1}{5}, 1)$, $v_2 = (0, 1)$, $v_3 = (-\frac{1}{4}, 1)$ on the edges of the lower hull of the Newton polytope of the homotopy $h(x, t) = 0$.

The inner normal v_1 attains the minimal inner product with those vertices on the first edge of the lower hull. Consider the four values of the inner product of v_1 with the four vertices of the lower hull:

$$\langle \left(\frac{1}{5}, 1\right), \{(0,1),(5,0),(7,0),(11,1)\}\rangle = \left\{1, 1, \frac{7}{5}, \frac{16}{5}\right\}.$$

Indeed, the minimal values occur with the first two vertices which span the first edge. This geometric construction motivates the following change of coordinates: let $x = yt^{1/5}$, we obtain

$$h(y,t) = t - y^5 t + y^7 t^{7/5} - y^{11} t^{16/5} \tag{8.8}$$

$$= t\left(1 - y^5 + y^7 t^{2/5} - y^{11} t^{11/5}\right). \tag{8.9}$$

We see that $\frac{1}{t}h(y,0) = 1 - y^5 = 0$ has one positive real root: $y = 1$. Now we can choose $t = \Delta t > 0$ such that Newton's method converges quadratically to a positive real root starting at $y = 1$. Let y^*: $h(y^*, \Delta t) = 0$, then we find the corresponding root in the original coordinates as $x^* = y^*(\Delta t)^{1/5}$.

We can even explicitly construct the fractional power series using Newton's method in a computer algebra system like Maple. The following sequence of Maple commands achieve this:

```
> h := t-x^5 + x^7 - x^(11)*t:
> hy := subs(x = y*t^(1/5),h):
> hyt := simplify(hy/t):
> newton := x -> x - subs(y=x,hyt/diff(hyt,y)):
> x[0] := 1:
> for k from 1 to 6 do
>   x[k] := newton(x[k-1]):
>   s[k] := series(x[k],t=0,15):
>   lprint(op(1,s[k]-s[k-1]));
> end do:
```

The output of the loop (done in Maple 9) shows the errors between two consecutive series expansions:

```
1
-301/15625*t^2
-84/3125*t^2
-2112/1953125*t^(18/5)
-32768/152587890625*t^(32/5)
-2147483648/23283064365386962890625*t^(64/5)
```

We observe the quadratic convergence, typical for Newton's method. While the particular values for the errors shows above may differ on other platforms with different versions of Maple, the computed fractional power series expansion is "exact", here we see the series up to third order:

```
> series(x[6],t=0,3);

       2/5    4/5       6/5        8/5          2    11/5
      t      t      34 t       266 t       11284 t    t
  1 + ---- + ---- + ------- + -------- + -------- - -----
       5      5       125        625        15625       5

                   12/5       13/5       14/5
            100947 t       14 t       12 t          3
          + ------------ - -------- + -------- + O(t )
              78125          25         5
```

To find the third positive real root, we proceed in a similar fashion, using the third inner normal $v_3 = (-1/4, 1)$ in the coordinate change $x = yt^{-1/4}$. As it turns out, we can take Δt quite large. For $\Delta t = 0.1$, $h(x, 0.1) = 0$ has the following three positive (approximate) real roots: 0.73, 1.0, and 1.56. As Δt grows larger, the real roots collide into multiple roots before escaping to the complex plane.

Exercise 8.2.6. Compute the fractional power series for the third positive real root, using Newton's method like shown above. Make sure enough terms in the series expansions are used so that the quadratic convergence is obvious.

In numerical implementations of polyhedral homotopies, we only use the first term of the fractional power series (also known as Puiseux series). The connection between these fractional power series and Newton polygons is classical for polynomials in two variables, see for example [Lef53] or [Wal62]. The generalization to systems of equations can be found in [McD02].

Using Newton polytopes to construct real curves and hypersurfaces with a prescribed topology is done by Viro's method [IS03, IV96]. This homotopy to glue real roots can be generalized to the case of complete intersections by the use of mixed subdivisions, see [Stu94b, Stu94c]. We will define these mixed subdivisions in the next section. We apply these co-called polyhedral homotopies to solve generic polynomial systems with given fixed Newton polytopes.

8.2.5 The Cayley trick and Minkowski's theorem

Mixed volumes were defined by Minkowski who showed that the volume of a linear combination of polytopes is a homogeneous polynomial in the factors of of the combination. The coefficients of this polynomial are mixed volumes. We will visualize this theorem on a simple example by the Cayley trick.

The Cayley trick [GKZ94, Proposition 1.7, page 274] is a method to rewrite a certain resultant as a discriminant of one single polynomial with additional variables. The polyhedral version of this trick as in [Stu94a, Lem. 5.2] is due to Bernd Sturmfels. See [HRS00] for another application of this trick.

Consider the following system:

$$\begin{array}{ll} f = (f_1, f_2) & \mathcal{A} = (A_1, A_2) \\ = \begin{cases} x_1^3 x_2 + x_1 x_2^2 + 1 = 0 \\ x_1^4 + x_1 x_2 + 1 = 0 \end{cases} & \begin{array}{l} A_1 = \{(3,1),(1,2),(0,0)\} \\ A_2 = \{(4,0),(1,1),(0,0)\} \end{array} \end{array}$$

The sparse structure of f is modeled by the tuple $\mathcal{A} = (A_1, A_2)$, where A_1 and A_2 are the *supports* of f_1 and f_2 respectively. The Newton polytopes are the convex hulls of the supports. The Cayley polytope of r polytopes is the convex hull of the polytopes placed at the vertices of an $(r-1)$-dimensional unit simplex. Figure 8.7 illustrates this construction for our example.

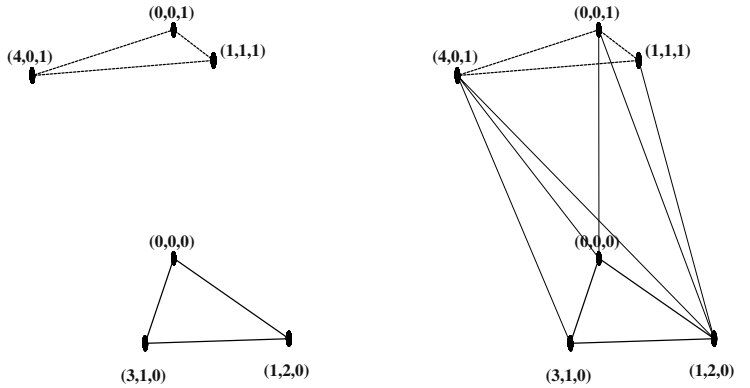

Fig. 8.7. The Cayley polytope of two polygons. The first polygon is placed at the vertex $(0,0,0)$, the second polygon is placed at $(0,0,1)$.

For our example, the Cayley polytope is so simple that a triangulation is obvious (see Figure 8.8). As every simplex has four vertices, either the simplex has three vertices from the same polygon (and the fourth one of the other polygon), or the simplex has two vertices of each polygon. A simplex of the first type is called unmixed, a simplex of the second type is mixed. Imagine taking slices parallel to the base of the Cayley polytope. These slices produce scaled copies of the original polygons in the unmixed simplices. In

the mixed simplex we find one scaled edge from the first and another scaled edge from the second polygon, see Figure 8.8.

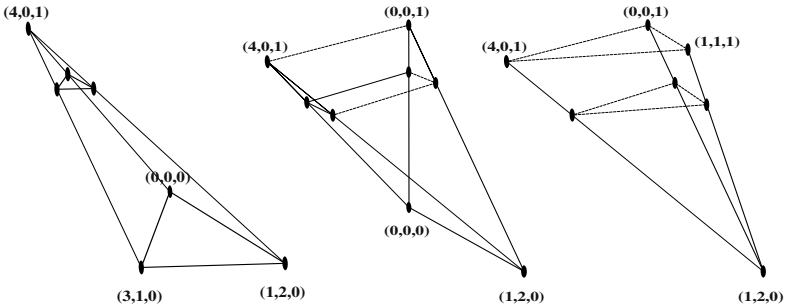

Fig. 8.8. A triangulation of the Cayley polytope. The middle simplex is mixed, the other two simplices are unmixed.

On Figure 8.9 we see in the cross section of the Cayley polytope a mixed subdivision of the convex combination $\lambda_1 P_1 + \lambda_2 P_2$, $\lambda_1 + \lambda_2 = 1$, $\lambda_1 \geq 0$ and $\lambda_2 \geq 0$, where P_1 defines the base and P_2 is at the top of the polytope. The areas of the triangles in the cross section are $\lambda_1^2 \times \text{area}(P_1)$ and $\lambda_2^2 \times \text{area}(P_2)$, as each side of the triangle is scaled by λ_1 and λ_2 respectively. The area of the cell in the subdivision spanned by one edge of P_1 (scaled by λ_1) and the other edge of P_2 (scaled by λ_2) is scaled by $\lambda_1 \times \lambda_2$, as we move the cross section.

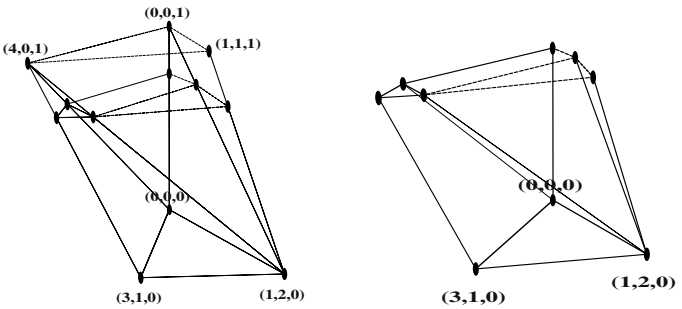

Fig. 8.9. A mixed subdivision induced by a triangulation of the Cayley polytope.

In Figure 8.10 we show the Minkowski sum of the two polygons P_1 and P_2, with their mixed subdivision corresponding to the triangulation of the Cayley polytope. For this example, Minkowski's theorem becomes

$$\begin{aligned}\text{area}(\lambda_1 P_1 + \lambda_2 P_2) &= V(P_1, P_1)\lambda_1^2 + V(P_1, P_2)\lambda_1\lambda_2 + V(P_2, P_2)\lambda_2^2 \\ &= 3\lambda_1^2 + 8\lambda_1\lambda_2 + 2\lambda_2^2.\end{aligned} \tag{8.10}$$

The coefficients in the polynomial (8.10) are mixed volumes (or areas in our example): $V(P_1, P_1)$ and $V(P_2, P_2)$ are the respective areas of P_1 and P_2, while $V(P_1, P_2)$ is the mixed area.

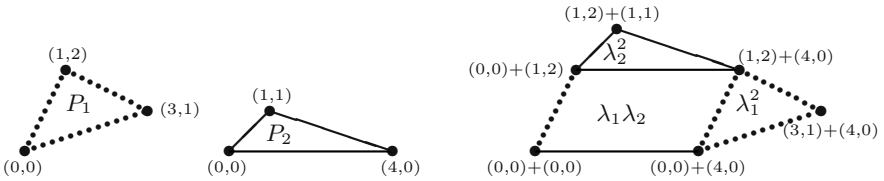

Fig. 8.10. A subdivision of the sum of two polygons P_1 and P_2. The sum is the convex hull of all sums of the vertices of the polygons. The cells in the subdivision are labeled by the multipliers for the area of $\lambda_1 P_1 + \lambda_2 P_2$.

The subdivisions we need are induced by a lifting. Such subdivisions are called regular, they define polyhedral homotopies. For the example, the lifted supports are $\widehat{\mathcal{A}} = (\widehat{A}_1, \widehat{A}_2)$, with

$$\widehat{A}_1 = \{(3,1,1), (1,2,0), (0,0,0)\} \text{ and } \widehat{A}_2 = \{(4,0,0), (1,1,1), (0,0,0)\}.$$

Figure 8.11 shows the mixed subdivision of Figure 8.10 as induced by the lower hull of the sum of the lifted polytopes.

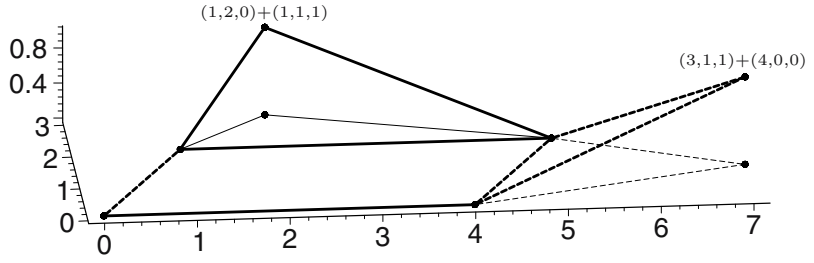

Fig. 8.11. A mixed subdivision is regular if it is induced by a lifting.

As there is only one mixed cell in the mixed subdivision of the Newton polytopes of our example, there is only one homotopy to consider, for example:

$$h(\mathbf{x}, t) = \begin{cases} x_1^3 x_2 t + x_1 x_2^2 + 1 = 0 \\ x_1^4 + x_1 x_2 t + 1 = 0 \end{cases} \tag{8.11}$$

The powers of the t in $h(\mathbf{x}, t) = \mathbf{0}$ are the lifting values of the supports which induced the mixed subdivision shown in Figure 8.11.

Exercise 8.2.7. Verify that the start system $h(\mathbf{x}, t = 0) = \mathbf{0}$ in the polyhedral homotopy (8.11) has indeed eight ($= V(P_1, P_2)$) regular solutions. Show that any system with exactly two monomials in every equation has always as many regular roots as its mixed volume, for any nonzero choice of the coefficients.

8.2.6 Computing mixed volumes and polyhedral continuation

In the previous subsections we introduced polyhedral continuation and mixed volumes. With these two concepts we can state and prove Bernshteĭn's first theorem. As the way we compute mixed volumes determines the way we solve a generic system, this section presents two different methods to compute mixed volumes. The first technique relies on the Cayley trick and computes all cells in a mixed subdivision. The second method uses linear programming and leads to an efficient enumeration of all mixed cells in a mixed subdivision.

With the Cayley trick we can obtain a regular mixed subdivision as a regular triangulation of the Cayley polytope. We next introduce a method to compute a regular triangulation of any polytope. Our method will construct the triangulation incrementally, adding the points one after the other. The key operation is to decompose one point with respect to one simplex. Consider for example the simplex $[\mathbf{c}_0, \mathbf{c}_1, \mathbf{c}_2]$ spanned by $\mathbf{c}_0 = (0, 0)$, $\mathbf{c}_1 = (3, 2)$, and $\mathbf{c}_2 = (2, 4)$. If we take one extra point, three possible updates can occur, illustrated by Table 8.2.

point	barycentric decomposition	pivoting
$\mathbf{x} = (2, 3)$:	$\mathbf{x} = +\frac{1}{8}\,\mathbf{c}_0 + \frac{1}{4}\,\mathbf{c}_1 + \frac{5}{8}\,\mathbf{c}_2$	no new simplex
$\mathbf{y} = (5, 1)$:	$\mathbf{y} = -\frac{1}{3}\,\mathbf{c}_0 + \frac{9}{4}\,\mathbf{c}_1 - \frac{7}{8}\,\mathbf{c}_2$	$[\mathbf{y}, \mathbf{c}_1, \mathbf{c}_2][\mathbf{c}_0, \mathbf{c}_1, \mathbf{y}]$
$\mathbf{z} = (1, 5)$:	$\mathbf{z} = +\frac{1}{8}\,\mathbf{c}_0 - \frac{3}{4}\,\mathbf{c}_1 + \frac{13}{8}\,\mathbf{c}_2$	$[\mathbf{c}_0, \mathbf{z}, \mathbf{c}_2]$

Table 8.2. Three possible updates of the simplex $[\mathbf{c}_0, \mathbf{c}_1, \mathbf{c}_2]$ with one point, $\mathbf{x}$, $\mathbf{y}$, or $\mathbf{z}$. Either we have no, two, or one new simplex by interchanging the vertex with negative coefficient with the point.

Solving a linear system we can write any point as a linear combination of the vertices of a simplex, requiring the coefficients in that linear combination to sum up to one. We call this linear combination a barycentric decomposition of a point with respect to a simplex. The negative signs of the coefficients in this barycentric decomposition tell which vertices of the simplex to interchange with the new point to create new simplices in the triangulation of the convex hull of the original simplex and the point. As we can see from Fig-

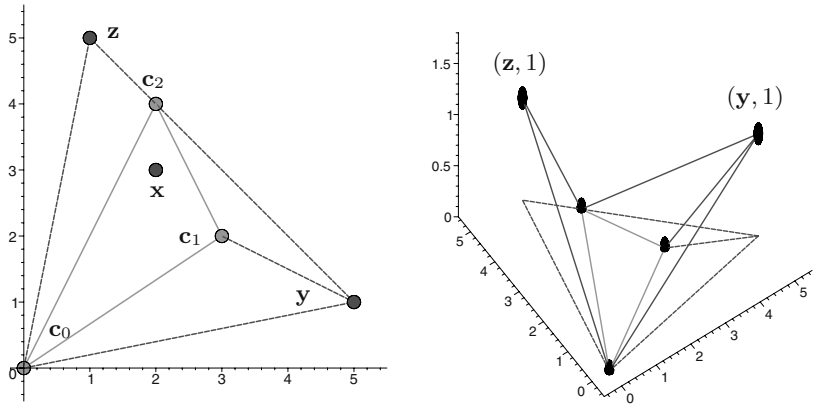

Fig. 8.12. Pivoting to obtain a regular triangulation of a polygon. The construction on the right shows how the triangulation can be obtained as the lower hull of $\mathbf{y}$ and $\mathbf{z}$ lifted at height one, with $[\mathbf{c}_0, \mathbf{c}_1, \mathbf{c}_2]$ sitting at level zero.

ure 8.12, any triangulation obtained by placing points (see [Lee91] for more on triangulations) in this way is regular.

The algorithm to compute regular triangulations incrementally leads to an incremental polyhedral solver, which solves polynomial systems adding one monomial after the other, see [VGC96]. If the structure of a polynomial system is such that most polynomials share the same support (or more generally span the same Newton polytope), and thus there are only few distinct Newton polytopes to consider, then the Cayley trick is not too wasteful.

The complexity of computing volumes and mixed volumes is discussed respectively in [DF88] and [DGH98].

Theorem 8.2.8. (Bernshteĭn's theorem A) *The number of roots of a generic system equals the mixed volume of its Newton polytopes.*

In his proof of this theorem, Bernshteĭn [Ber75] used a homotopy (implemented in [VVC94]), based on a recursive formula for computing mixed volumes. This proof idea was generalized by Huber and Sturmfels in [HS95]. Note that the theorem concerns "generic systems", which are systems with randomly chosen coefficients. These generic systems serve as start system in a coefficient-parameter homotopy to solve any specific polynomial system with the same Newton polytopes.

For the coordinate changes in the polyhedral homotopies, we need to know the inner normals to the mixed cells. Therefore, we use a dual representation of polytopes, see Figure 8.13. The normal fan of a polytope is the collection of the normal cones to all faces of the polytope. The normal cone to a face contains all inner normals which define the face.

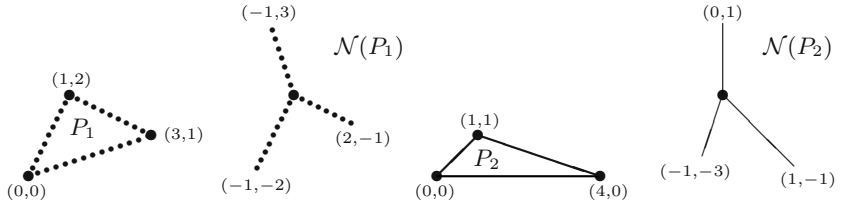

Fig. 8.13. Two polygons P_1 and P_2 and their normal fans, $\mathcal{N}(P_1)$ and $\mathcal{N}(P_2)$. The labels corresponding to the edges in the fans are inner normals to the corresponding edges of the polygons.

We are only interested in the mixed cells of a mixed subdivision, and in particular, the inner normal to those lower facets of the Minkowski sum which define the mixed cells. Figure 8.14 illustrates that the inner normal to a mixed cell lies in the intersection of the normal cones to the edges which span that mixed cell.

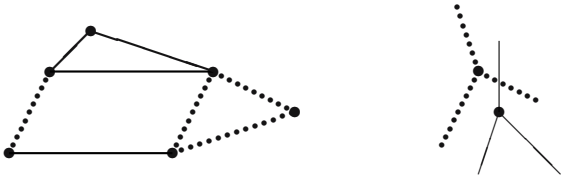

Fig. 8.14. The dual representation of a mixed subdivision.

The search for all inner normals to the mixed cells in a mixed subdivision naturally leads to a system of linear equalities and inequalities. For a tuple of n supports $(A_1, A_2, \ldots, A_n)$, consider an edge of the kth polytope, spanned by $\{\mathbf{a}, \mathbf{b}\} \subseteq A_k$. Then the inner normal $\mathbf{v}$ to this edge satisfies

$$\begin{cases} \langle \mathbf{a}, \mathbf{v} \rangle = \langle \mathbf{b}, \mathbf{v} \rangle \\ \langle \mathbf{a}, \mathbf{v} \rangle \leq \langle \mathbf{c}, \mathbf{v} \rangle, \quad \text{for all } \mathbf{c} \in A_k. \end{cases} \tag{8.12}$$

Enumerating all edges of a polytope is thus equivalent to enumerating all feasible solutions to the system (8.12). Letting k range from 1 to n in (8.12) applied to the lifted point sets $\widehat{A}_k$ provides the dual linear-programming model to enumerate all inner normals to the mixed cells in a regular mixed subdivision.

A lift-and-prune strategy to enumerate all mixed cells in a regular mixed subdivision was proposed in [EC95] and dualized in [VGC96]. Recently, insight in the linear programming methods has led to very efficient calculations of mixed volumes, as developed in [DKK03], [GL00, GL03], [KK03b], [LL01], and [TKF02].

8.2.7 Bernshteĭn's second theorem

When tracing solution paths diverging to infinity, one may wonder when to stop. After all, infinity is pretty far off, and even if good knowledge of the application domain gives us good bounds on the size of the solutions, we do not want to miss valid solutions with large components. If a path seems to diverge, we must know whether we have true divergence or convergence to a root with large components. Bernshteĭn's second theorem [Ber75] will provide us with a certificate of divergence.

For a system $f(\mathbf{x}) = 0$, supported by $\mathcal{A} = (A_1, A_2, \ldots, A_n)$, we can write its equations $f = (f_1, f_2, \ldots, f_n)$ as

$$f_i(\mathbf{x}) = \sum_{\mathbf{a} \in A_i} c_{i\mathbf{a}} \mathbf{x}^{\mathbf{a}}, \quad i = 1, 2, \ldots, n.$$

The Newton polytopes of f are denoted by $\mathcal{P} = (P_1, P_2, \ldots, P_n)$, with $P_i := \text{conv}(A_i)$, $i = 1, 2, \ldots, n$. Then for any $\omega \neq \mathbf{0}$, we define the tuple of faces $\partial_\omega \mathcal{P} = (\partial_\omega P_1, \partial_\omega P_2, \ldots, \partial_\omega P_n)$, as $\partial_\omega P_i := \text{conv}(\partial_\omega A_i)$, with

$$\partial_\omega A_i := \{\ \mathbf{a} \in A_i \mid \langle \mathbf{a}, \omega \rangle = \min_{\mathbf{a}' \in A_i} \langle \mathbf{a}', \omega \rangle\ \}. \tag{8.13}$$

The set $\partial_\omega A_i$ is the support of the face of the ith polynomial f_i:

$$\partial_\omega f_i(\mathbf{x}) = \sum_{\mathbf{a} \in \partial_\omega A_i} c_{i\mathbf{a}} \mathbf{x}^{\mathbf{a}}.$$

We write $\partial_\omega f = (\partial_\omega f_1, \partial_\omega f_2, \ldots, \partial_\omega f_n)$ as the face of the system f determined by $\omega \neq \mathbf{0}$. The mixed volume of $\mathcal{P}$ is denoted by $V(\mathcal{P})$ and $\mathbb{C}^* = \mathbb{C} \setminus \{0\}$.

Theorem 8.2.9. (Bernshteĭn's theorem B) *If $\forall \omega \neq \mathbf{0}$, $\partial_\omega f(\mathbf{x}) = \mathbf{0}$ has no solutions in $(\mathbb{C}^*)^n$, then $V(\mathcal{P})$ is exact and all solutions are isolated. Otherwise, for $V(\mathcal{P}) \neq 0$: $V(\mathcal{P}) >$ #isolated solutions.*

Interestingly, the Newton polytopes may often be in general position, i.e.: $V(\mathcal{P})$ is exact for every nonzero choice of the coefficients. Consider for example the following system:

$$f(\mathbf{x}) = \begin{cases} c_{111}x_1x_2 + c_{110}x_1 + c_{101}x_2 + c_{100} = 0 \\ c_{222}x_1^2x_2^2 + c_{210}x_1 + c_{201}x_2 = 0 \end{cases}$$

We show the tuple of Newton polytopes in Figure 8.15.

Exercise 8.2.10. Verify that the mixed volume $V(P_1, P_2)$ of the polygons P_1 and P_2 is indeed equal to four.

While the observation in Figure 8.15 would let us believe that the mixed volume always provides a sharp root count, we have to keep in mind that the vertices of the polytopes are not randomly chosen. The vertices occur as

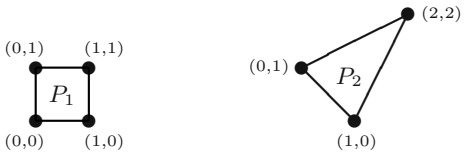

Fig. 8.15. Two Newton polygons in general position: $\forall \omega \neq \mathbf{0} : \partial_\omega A_1 + \partial_\omega A_2 \leq 3 \Rightarrow V(P_1, P_2) = 4$ is always exact, for all nonzero choices of the coefficients of f, because we need at least four monomials for $\partial_\omega f(\mathbf{x}) = \mathbf{0}$ to have all its roots in $(\mathbb{C}^*)^2$.

the exponents in the polynomials. For instance, general Newton polytopes are almost never simplicial, we usually find k-dimensional faces spanned by far more than $k+1$ vertices.

Following Bernshteĭn we look at what happens when we consider the solution paths in a homotopy going from a generic to a specific polynomial system. At the limit of the paths, we look at the power series expansion, using the following result.

Theorem 8.2.11. $\forall \mathbf{x}(t)$, $h(\mathbf{x}(t), t) = (1-t)g(\mathbf{x}(t)) + tf(\mathbf{x}(t)) = \mathbf{0}$,
$\exists s > 0$, $m \in \mathbb{N} \setminus \{0\}$, $\omega \in \mathbb{Z}^n$: $\begin{cases} x_i(s) = b_i s^{\omega_i}(1 + O(s)), & i = 1, 2, \ldots, n \\ t(s) = 1 - s^m & \text{for } t \approx 1, s \approx 0 \end{cases}$

The number m is called the *winding number* of the solution at the end of the path (not to be confused with the multiplicity). The winding number is the smallest number so that $\mathbf{z}(2\pi m) = \mathbf{z}(0)$, if we consider $\mathbf{z}(\theta)$ a solution path of $h(\mathbf{z}(\theta), t(\theta)) = \mathbf{0}$, winding around 1 with values for the continuation parameter t defined by $t = 1 + (t_0 - 1)e^{i\theta}$, as $t_0 \approx 1$.

At the end of a path, when does $\lim_{t \to 1} x_i(t) \in \mathbb{C}^*$? From Theorem 8.2.11, we can characterize the divergence of the path $\mathbf{x}(t)$ by the leading exponents ω in the power series:

$$x_i(t) \begin{cases} \to \infty \\ \in \mathbb{C}^* \\ \to 0 \end{cases} \quad \Leftrightarrow \quad \omega_i \begin{cases} < 0 \\ = 0 \\ > 0 \end{cases}$$

From this simple observation we see that a solution at infinity and a solution with zero components are regarded (or disregarded) equally.

Next we show the relation between face systems and power series. Assuming $\lim_{t \to 1} x_i(t) \notin \mathbb{C}^*$, and $\omega_i \neq 0$, we consider a diverging path.

First we substitute the power series $x_i(s) = b_i s^{\omega_i}(1 + O(s))$, $i = 1, 2, \ldots, n$, $t(s) = 1 - s^m$, $s \approx 0$ into the homotopy $h(\mathbf{x}, t) = (1-t)g(\mathbf{x}) + tf(\mathbf{x}) = \mathbf{0}$. We find

$$h(\mathbf{x}(s), t(s)) = \underbrace{f(\mathbf{x}(s))}_{\text{dominant as } s \to 0} + s^m(g(\mathbf{x}(s)) - f(\mathbf{x}(s))) = \mathbf{0}.$$

Thus (as expected), the choice of the start system $g(\mathbf{x}) = \mathbf{0}$ plays no role in what happens as s approaches zero. Let us now see what the substitution does to the ith polynomial:

$$f_i(\mathbf{x}) = \sum_{\mathbf{a}\in A_i} c_{i\mathbf{a}}\mathbf{x}^{\mathbf{a}} \to f_i(\mathbf{x}(s)) = \underbrace{\sum_{\mathbf{a}\in A_i} c_{i\mathbf{a}} \prod_{i=1}^{n} b_i^{a_i} s^{\langle \mathbf{a},\omega\rangle}}_{\partial_\omega f_i(\mathbf{x}(s)) \text{ dominant}} (1+O(s)).$$

Arranging the monomials in $f(\mathbf{x}(s))$ in increasing order of powers of s, we see that the monomials that become dominant as $s \to 0$ have exponents whose inner product is minimal with ω. Recall that we characterize these exponents by the face of the support A_i in the direction of ω, see (8.13). Moreover, as $f_i(\mathbf{x}(s)) = 0$ for $s \to 0$, we see from the result of the substitution that then $\partial_\omega f_i(\mathbf{b}) = 0$, and thus $\partial_\omega f(\mathbf{b}) = \mathbf{0}$ for some $\mathbf{b} \in (\mathbb{C}^*)^n$.

This is the key idea in the proof of Bernshteĭn's second theorem. Like his first theorem, his idea is very constructive: follow the direction of a diverging path and (in addition to a solution at infinity) we find a face system which has solutions in $(\mathbb{C}^*)^n$. This face system forms a certificate for the mixed volume to overshoot the actual number of roots.

That Richardson extrapolation is useful to find ω is not so surprising. A closer inspection of the errors of the error expansion reveals that a similar extrapolation scheme can be applied to approximate the winding number m.

As we get closer to our target system, we have to decrease our step size when dealing with a difficult path. For the purpose of extrapolation, we better decrease the step size geometrically, i.e., for some λ, $0 < \lambda < 1$, consecutive values $t_0, t_1, \ldots t_k$ of the continuation parameter t satisfy $1 - t_k = \lambda(1 - t_k) = \cdots = \lambda^k(1 - t_0)$ and for the corresponding sequence of s-values we have $s_k = \lambda^{1/m} s_{k-1} = \cdots = \lambda^{k/m} s_0$.

Recall the form of the power series for a solution path $\mathbf{x}(s)$ for s approaching zero: $x_i(s) = b_i s^{\omega_i}(1 + O(s))$ with $t(s) = 1 - s^m$. Sampled along $s_0, s_1, \ldots, s_k$, we obtain

$$x_i(s_k) = b_i \lambda^{k\omega_i/m} s_0^{\omega_i}(1 + O(\lambda^{k/m} s_0)). \tag{8.14}$$

Since we are interested in the leading powers ω_i, we take the logarithms of the magnitudes of the points sampled along the path:

$$\log|x_i(s_k)| = \log|b_i| + \frac{k\omega_i}{m}\log(\lambda) + \omega_i \log(s_0) + \log\left|1 + \sum_{j=0}^{\infty} b'_j (\lambda^{k/m} s_0)^j\right|.$$

A first-order approximation for ω_i is given by v_{kk+1} with the general extrapolation formula in $v_{k..l}$:

$$v_{kk+1} := \log|x_i(s_k + 1)| - \log|x_i(s_k)|, \quad v_{k..l} = v_{k..l-1} + \frac{v_{k+1..l} - v_{k..l-1}}{1 - \lambda}$$

which results in $\omega_i = m\frac{v_{0..r}}{\log(\lambda)} + O(s_0^r)$. While we can make the order r of the extrapolation as high as we like (thereby increasing the accuracy of ω_i). Notice that the formula assumes we know the winding number m.

If we examine the expansion of the errors:

$$e_i^{(k)} = (\log|x_i(s_k)| - \log|x_i(s_{k+1})|) \tag{8.15}$$
$$-(\log|x_i(s_{k+1})| - \log|x_i(s_{k+2})|) \tag{8.16}$$
$$= c_1\lambda^{k/m}s_0(1 + 0(\lambda^{k/m})), \tag{8.17}$$

we find similar extrapolation formulas to approximate m:

$$e_i^{(kk+1)} := \log(e_i^{(k+1)}) - \log(e_i^{(k)}), \quad e_i^{(k..l)} = e_i^{(k+1..l)} + \frac{e_i^{(k..l-1)} - e_i^{(k+1..l)}}{1 - \lambda_{k..l}}$$

with $\lambda_{k..l} = \lambda^{(l-k-1)/m_{k..l}}$. So we obtain $m_{k..l} = \frac{\log(\lambda)}{e_i^{(k..l)}} + O(\lambda^{(l-k)k/m})$

The system of Cassou-Noguès is a very nice example. It illustrates how symbolic results can be obtained by purely numerical means.

$$f(b,c,d,e) = \begin{cases} 15b^4cd^2 + 6b^4c^3 + 21b^4c^2d - 144b^2c - 8b^2c^2e \\ -28b^2cde - 648b^2d + 36b^2d^2e + 9b^4d^3 - 120 = 0 \\ 30c^3b^4d - 32de^2c - 720db^2c - 24c^3b^2e - 432c^2b^2 + 576ec \\ -576de + 16cb^2d^2e + 16d^2e^2 + 16e^2c^2 + 9c^4b^4 + 5184 \\ +39d^2b^4c^2 + 18d^3b^4c - 432d^2b^2 + 24d^3b^2e - 16c^2b^2de - 240c = 0 \\ 216db^2c - 162d^2b^2 - 81c^2b^2 + 5184 + 1008ec - 1008de \\ +15c^2b^2de - 15c^3b^2e - 80de^2c + 40d^2e^2 + 40e^2c^2 = 0 \\ 261 + 4db^2c - 3d^2b^2 - 4c^2b^2 + 22ec - 22de = 0 \end{cases}$$

Root counts: $D = 1344$, $B = 312$, $V(\mathcal{P}) = 24$, but there are only 16 finite roots.

$$\partial_{(0,0,0,-1)}f(b,c,d,e) = \begin{cases} -8b^2c^2e - 28b^2cde + 36b^2d^2e = 0 \\ -32de^2c + 16d^2e^2 + 16e^2c^2 = 0 \\ -80de^2c + 40d^2e^2 + 40e^2c^2 = 0 \\ 22ec - 22de = 0 \end{cases}$$

The winding number is $m = 2$. See [HV98] for more about polyhedral end games.

8.3 Homotopies for positive dimensional solution sets

To introduce the numerical representation of positive dimensional solution sets, we start off with a dictionary, linking concepts in algebraic geometry to data and algorithms in numerical analysis. Witness sets form the central

data and are obtained by a cascade of homotopies. The companion algorithms to the witness sets are membership tests to decide whether any given point belongs to a certain component of the solution set. We illustrate a numerical irreducible decomposition on a simple example and give an overview of our numerical factorization methods.

8.3.1 A dictionary

Kempf writes in [Kem93] that "Algebraic geometry studies the delicate balance between the geometrically plausible and the algebraically possible". With our numerical tools, we feel closer to the geometrical than to the algebraic side, because we are not calculating with polynomials in the algebraic sense. In [SVW03] we outlined the structure of a dictionary, presented as Table 8.3.

Numerical Algebraic Geometry Dictionary		
Algebraic Geometry	example in 3-space	Numerical Analysis
variety	collection of points, algebraic curves, and algebraic surfaces	polynomial system + union of witness sets, see below for the definition of a witness point
irreducible variety	a single point, or a single curve, or a single surface	polynomial system + witness set + probability-one membership test
generic point on an irreducible variety	random point on an algebraic curve or surface	point in a witness set; a witness point is a solution of the polynomial system on the variety and on a random slice whose codimension is the dimension of the variety
pure dimensional variety	one or more points, or one or more curves, or one or more surfaces	polynomial system + set of witness sets of same dimension + probability-one membership tests
irreducible decomposition of a variety	several pieces of different dimensions	polynomial system + array of sets of witness sets and probability-one membership tests

Table 8.3. Dictionary to translate algebraic geometry into numerical analysis.

8.3.2 Witness sets and a cascade of homotopies

A witness set is the basic concept of numerical algebraic geometry as it allows us to apply numerical methods for isolated solutions to positive dimensional solution components.

Every irreducible component of a solution set is presented by a witness set whose cardinality equals the degree of the irreducible component. To reduce

a solution set of dimension k to a set of isolated points, we cut the k degrees of freedom by adding k random hyperplanes $L(\mathbf{x}) = \mathbf{0}$ to the system $f(\mathbf{x}) = \mathbf{0}$ which defines the entire solution set.

One obstacle is that we have to deal with systems whose number of equations in not necessarily the same as the number of unknowns. If there are fewer equations than unknowns, we simply add enough random hyperplanes to make up for the difference, so underdetermined systems are easy to handle.

Let us consider overdetermined systems, say f consists of 5 equations in 3 variables. To turn f into a system of N equations in N variables where N is either 3 or 5, we can respectively apply the following techniques:

randomization: Choosing random complex numbers a_{ij}, we add random combinations of the last two polynomials to the first three polynomials:

$$\begin{cases} f_1(\mathbf{x}) + a_{11} f_4(\mathbf{x}) + a_{12} f_5(\mathbf{x}) = 0 \\ f_2(\mathbf{x}) + a_{21} f_4(\mathbf{x}) + a_{22} f_5(\mathbf{x}) = 0 \\ f_3(\mathbf{x}) + a_{31} f_4(\mathbf{x}) + a_{32} f_5(\mathbf{x}) = 0 \end{cases}$$

slack variables: We introduce two new variables z_1 and z_2 (so-called slack variables) and add random multiples of these variables to every equation:

$$\begin{cases} f_1(\mathbf{x}) + a_{11} z_1 + a_{12} z_2 = 0 \\ f_2(\mathbf{x}) + a_{21} z_1 + a_{22} z_2 = 0 \\ f_3(\mathbf{x}) + a_{31} z_1 + a_{32} z_2 = 0 \\ f_4(\mathbf{x}) + a_{41} z_1 + a_{42} z_2 = 0 \\ f_5(\mathbf{x}) + a_{51} z_1 + a_{52} z_2 = 0 \end{cases}$$

While the randomization technique might seem at first more attractive because we are left with fewer equations, working with slack variables provides a cascade of homotopies to compute candidate witness points on all positive dimensional components.

In particular, considering f_4 and f_5 as hyperplanes L_1 and L_2 to cut the solution set of the first three equations in f, we consider a cascade of three systems. To get witness points on the two dimensional solution sets, we first solve

$$\begin{cases} f_1(\mathbf{x}) + a_{11} z_1 + a_{12} z_2 = 0 \\ f_2(\mathbf{x}) + a_{21} z_1 + a_{22} z_2 = 0 \\ f_3(\mathbf{x}) + a_{31} z_1 + a_{32} z_2 = 0 \\ L_1(\mathbf{x}) + z_1 = 0 \\ L_2(\mathbf{x}) + z_2 = 0 \end{cases}$$

Solutions with $z_1 = 0$ and $z_2 = 0$ define witness points on the two dimensional solution components. Solutions with $z_1 \neq 0$ and $z_2 \neq 0$ provide start points in the homotopy which removes L_2 from the system, which leads to the next system in the cascade:

$$\begin{cases} f_1(\mathbf{x}) + a_{11}z_1 + a_{12}z_2 = 0 \\ f_2(\mathbf{x}) + a_{21}z_1 + a_{22}z_2 = 0 \\ f_3(\mathbf{x}) + a_{31}z_1 + a_{32}z_2 = 0 \\ L_1(\mathbf{x}) + z_1 = 0 \\ z_2 = 0 \end{cases}$$

The paths defined by this move end at witness points on the one dimensional components, picked out by $z_1 = 0$. Solutions with $z_1 \neq 0$ are used in the homotopy which removes L_1 to lead to the isolated solutions of the system. The last system in the cascade is

$$\begin{cases} f_1(\mathbf{x}) + a_{11}z_1 = 0 \\ f_2(\mathbf{x}) + a_{21}z_1 = 0 \\ f_3(\mathbf{x}) + a_{31}z_1 = 0 \\ z_1 = 0 \\ z_2 = 0 \end{cases}$$

In the next section we give a specific example of this cascade.

The idea of slicing a solution set by hyperplanes to determine its dimension appeared in [GH93] to prove that the theoretical complexity of this problem is polynomial.

Exercise 8.3.1. Consider the adjacent minors of a general 2×4-matrix:

$$\begin{bmatrix} x_{11} & x_{12} & x_{13} & x_{14} \\ x_{21} & x_{22} & x_{23} & x_{24} \end{bmatrix} \qquad f(\mathbf{x}) = \begin{cases} x_{11}x_{22} - x_{21}x_{12} = 0 \\ x_{12}x_{23} - x_{22}x_{13} = 0 \\ x_{13}x_{24} - x_{23}x_{14} = 0 \end{cases}$$

Verify that $\dim(f^{-1}(\mathbf{0})) = 5$ and $\deg(f^{-1}(\mathbf{0})) = 8$. This is the simplest instance of a general family of problems introduced in [DES98], see [HS00] for special decomposition methods.

8.3.3 A probability-one membership test

A probability-one membership test determines whether a given point $\mathbf{p}$ lies on a pure dimensional solution set. Suppose we have witness points defined by a polynomial system $f(\mathbf{x}) = \mathbf{0}$ and hyperplanes $L(\mathbf{x}) = \mathbf{0}$. A homotopy method implements the probability-one membership test:

1. Define $K(\mathbf{x}) = L(\mathbf{x}) - L(\mathbf{p})$. As $K(\mathbf{p}) = \mathbf{0}$, the hyperplanes K pass through $\mathbf{p}$.
2. Consider the homotopy

$$h(\mathbf{x}, t) = \begin{pmatrix} f(\mathbf{x}) \\ K(\mathbf{x}) \end{pmatrix} (1 - t) + \begin{pmatrix} f(\mathbf{x}) \\ L(\mathbf{x}) \end{pmatrix} t = \mathbf{0}.$$

At $t = 1$ we start tracking paths at the witness set and find their end points at $t = 0$.

3. If $\mathbf{p}$ belongs to the solution set of $h(\mathbf{x}, 0) = \mathbf{0}$, then it is also a witness point of the pure dimensional solution set.

Notice that this test does not move the point $\mathbf{p}$, which may be a highly singular point. This observation is important for the numerical stability of this test. The test is illustrated in Figure 8.16.

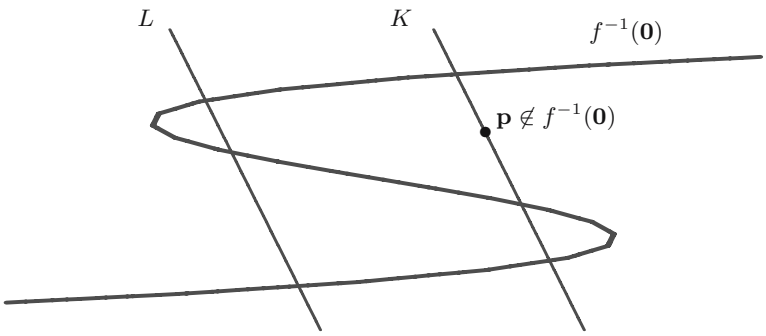

Fig. 8.16. Illustration of a probability-one membership test using a homotopy. The homotopy moves the line L of the witness set for $f^{-1}(0)$ to the line K, which passes to the test point $\mathbf{p}$. As none of the witness points on K equals $\mathbf{p}$, $\mathbf{p} \notin f^{-1}(0)$.

8.3.4 A numerical irreducible decomposition

Consider the following example:

$$f(\mathbf{x}) = \begin{cases} (x_1 - 1)(x_2 - x_1^2) = 0 \\ (x_1 - 1)(x_3 - x_1^3) = 0 \\ (x_1^2 - 1)(x_2 - x_1^2) = 0 \end{cases}$$

From its factored form we see that $f(\mathbf{x}) = \mathbf{0}$ has two solution components: the two dimensional plane $x_1 = 1$ and the twisted cubic $\{ (x_1, x_2, x_3) \mid x_2 - x_1^2 = 0,\ x_3 - x_1^3 = 0 \}$.

To describe the solution set of this system, we use a cascade of homotopies, the chart in Figure 8.17 illustrates the flow of data for this example.

Because the top dimensional component is of dimension two, we add two random hyperplanes to the system and make it square again by adding two slack variables z_1 and z_2:

$$e(\mathbf{x}, z_1, z_2) = \begin{cases} (x_1 - 1)(x_2 - x_1^2) + a_{11}z_1 + a_{12}z_2 = 0 \\ (x_1 - 1)(x_3 - x_1^3) + a_{21}z_1 + a_{22}z_2 = 0 \\ (x_1^2 - 1)(x_2 - x_1^2) + a_{31}z_1 + a_{32}z_2 = 0 \\ c_{10} + c_{11}x_1 + c_{12}x_2 + c_{13}x_3 + z_1 = 0 \\ c_{20} + c_{21}x_1 + c_{22}x_2 + c_{23}x_3 + z_2 = 0 \end{cases}$$

where all constants a_{ij}, $i = 1, 2, 3$, $j = 1, 2$, and c_{kl}, $k = 1, 2$, $l = 0, 1, 2, 3$ are randomly chosen complex numbers. Observe that when $z_1 = 0$ and $z_2 = 0$ the solutions to $e(\mathbf{x}, z_1, z_2) = \mathbf{0}$ satisfy $f(\mathbf{x}) = \mathbf{0}$. So if we solve $e(\mathbf{x}, z_1, z_2) = \mathbf{0}$ we will find a single witness point on the two dimensional solution component $x_1 = 1$ as a solution with $z_1 = 0$ and $z_2 = 0$. Using polyhedral homotopies, this requires the tracing of six solutions paths.

The embedding was proposed in [SV00] to find generic points on all positive dimensional solution components with a cascade of homotopies. In [SV00] it was proven that solutions with slack variables $z_i \neq 0$ are regular and, moreover, that those solutions can be used as start solutions in a homotopy to find witness points on lower dimensional solution components. At each stage of the algorithm, we call solutions with nonzero slack variables *nonsolutions*.

In the solution of $e(\mathbf{x}, z_1, z_2) = \mathbf{0}$, one path ended with $z_1 = 0 = z_2$, the five other paths ended in regular solutions with $z_1 \neq 0$ and $z_2 \neq 0$. These five "nonsolutions" are start solutions for the next stage, which uses the homotopy

$$h_2(\mathbf{x}, z_1, z_2, t) = \begin{cases} (x_1 - 1)(x_2 - x_1^2) + a_{11}z_1 + a_{12}z_2 = 0 \\ (x_1 - 1)(x_3 - x_1^3) + a_{21}z_1 + a_{22}z_2 = 0 \\ (x_1^2 - 1)(x_2 - x_1^2) + a_{31}z_1 + a_{32}z_2 = 0 \\ c_{10} + c_{11}x_1 + c_{12}x_2 + c_{13}x_3 + z_1 = 0 \\ z_2(1 - t) + (c_{20} + c_{21}x_1 + c_{22}x_2 + c_{23}x_3 + z_2)t = 0 \end{cases}$$

where t goes from one to zero, replacing the last hyperplane with $z_2 = 0$. Of the five paths, four of them converge to solutions with $z_1 = 0$. Of those four solutions, one of them is found to lie on the two dimensional solution component $x_1 = 1$, the other three are generic points on the twisted cubic. As there is one solution with $z_1 \neq 0$, we have one candidate left to use as a start point in the final stage, which searches for isolated solutions of $f(\mathbf{x}) = \mathbf{0}$. The homotopy for this stage is

$$h_1(\mathbf{x}, z_1, t) = \begin{cases} (x_1 - 1)(x_2 - x_1^2) + a_{11}z_1 = 0 \\ (x_1 - 1)(x_3 - x_1^3) + a_{21}z_1 = 0 \\ (x_1^2 - 1)(x_2 - x_1^2) + a_{31}z_1 = 0 \\ z_1(1 - t) + (c_{10} + c_{11}x_1 + c_{12}x_2 + c_{13}x_3 + z_1)t = 0 \end{cases}$$

which as t goes from 1 to 0, replaces the last hyperplane $z_1 = 0$. At $t = 0$, the solution is found to lie on the twisted cubic, so there are no isolated solutions.

The calculations are summarized in Figure 8.17. The breakup into irreducibles will be explained in the next section.

8.3.5 Factorization methods

A recent trend in computer algebra is the adaptation of symbolic methods to deal with approximate input data, which leads to the use of hybrid methods [CKW02]. One such problem is the factorization of multivariate polynomials, listed as a challenge in [Kal00]. Recent papers on this problem are [CGvH$^+$01, CGKW02], [GR01, GR02], [HWSZ00], and [Sas01].

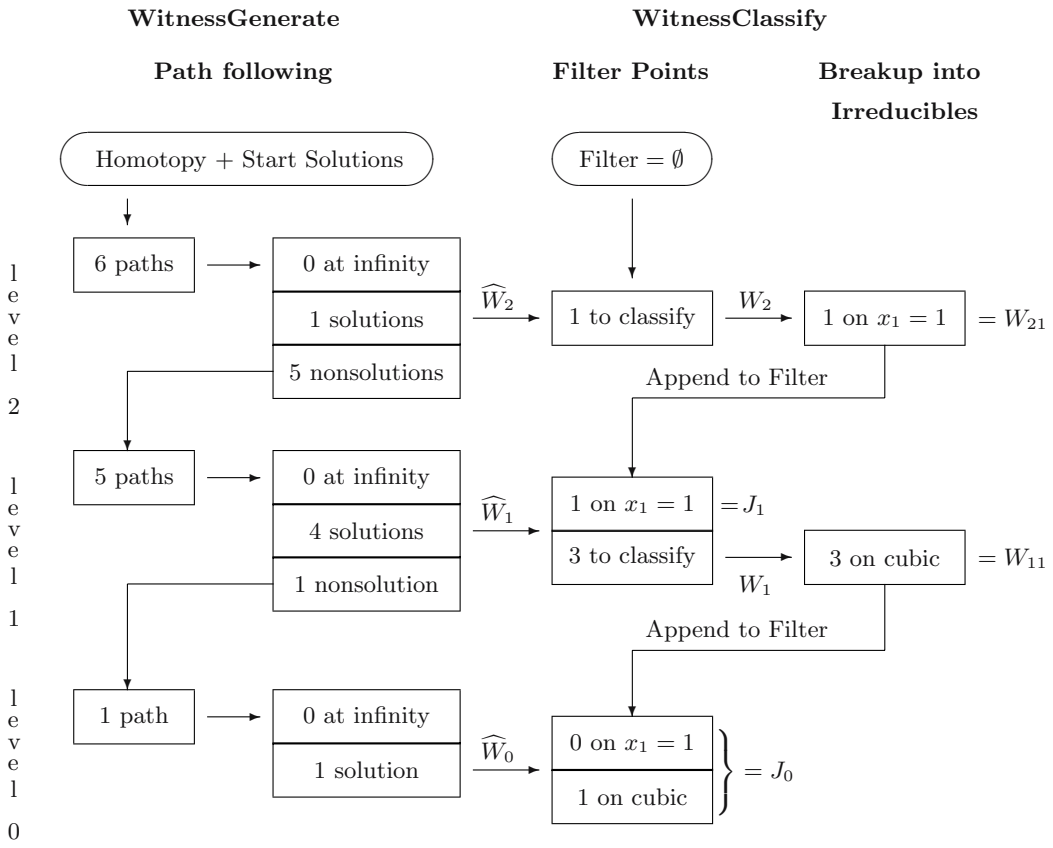

Fig. 8.17. Numerical Irreducible Decomposition of a system whose solutions are the 2-dimensional plane $x_1 = 1$ and the twisted cubic. At level i, for $i = 2, 1, 0$, we filter candidate witness sets $\widehat{W}_i$ into junk sets J_i and witness sets W_i. The sets W_i are partitioned into witness sets W_{ij} for the irreducible components.

Monodromy to partition witness point sets

We can see whether a curve factors or not by looking at its plot in complex space, i.e.: we consider the curve as a Riemann surface. Figure 8.18 was made with Maple (see [CJ98] for instructions).

Looking at Figure 8.18, imagine a line which intersects the surface in three points. Taking one complete turn of the line around the vertical axis $z = 0$ will cause the points to permute. For example, the point which was lowest will have moved up, while another point will have come down. Such a permutation can only happen if the corresponding algebraic curve is irreducible.

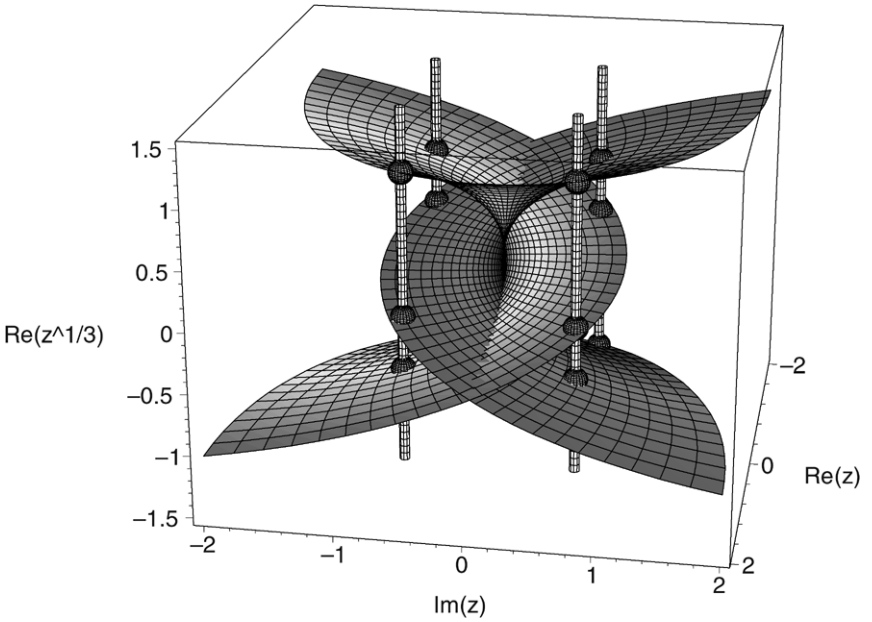

Fig. 8.18. The Riemann surface of $z^3 - w = 0$. The height of the surface is the real part of $w = z^{1/3}$, while the gray scale corresponds to the imaginary part of $w = z^{1/3}$. Observe that a loop around the origin permutes the order of points.

Based on this observation, we can decompose any pure dimensional set into irreducible components. Our monodromy algorithm returns a partition of the witness set for a pure dimensional component: points in the same subset of the partition belong to the same irreducible component. Recall that witness points are defined by a system $f(\mathbf{x}) = \mathbf{0}$ and a set of hyperplanes $L(\mathbf{x}) = \mathbf{0}$. With the homotopy

$$h_{KL}(\mathbf{x}, t) = \lambda \begin{pmatrix} f(\mathbf{x}) \\ K(\mathbf{x}) \end{pmatrix} (1 - t) + \begin{pmatrix} f(\mathbf{x}) \\ L(\mathbf{x}) \end{pmatrix} t = \mathbf{0}, \quad \lambda \in \mathbb{C},$$

we find new witness points on the hyperplanes $K(\mathbf{x}) = \mathbf{0}$, starting at those witness points satisfying $L(\mathbf{x}) = \mathbf{0}$, letting t move from one to zero. Choosing another random constant $\mu \neq \lambda$, we move back from K to L, using the homotopy

$$h_{LK}(\mathbf{x}, t) = \mu \begin{pmatrix} f(\mathbf{x}) \\ L(\mathbf{x}) \end{pmatrix} (1 - t) + \begin{pmatrix} f(\mathbf{x}) \\ K(\mathbf{x}) \end{pmatrix} t = \mathbf{0}, \quad \mu \in \mathbb{C}.$$

The homotopies $h_{KL}(\mathbf{x}, t) = \mathbf{0}$ and $h_{LK}(\mathbf{x}, t) = \mathbf{0}$ implement one loop in the monodromy algorithm, moving witness points from L to K and then back from K to L. At the end of the loop we have the same witness set as the set we started with, except possibly permuted. Permuted points belong to the same irreducible component.

Notice that the monodromy algorithm does not know the locations of the singularities. See [DvH01] for the algorithms to compute the monodromy group of an algebraic curve in Maple (package algcurves). Using homotopies theoretically, the complexity of factoring polynomials with rational coefficients was shown in [BCGW93] to be in NC.

Linear traces to validate the partition

When we run the monodromy algorithm, we may not have made enough loops to group as many witness points as the degree of each factor, i.e.: the partition predicted by the monodromy might be too fine. For a k-dimensional solution component, it suffices to consider a curve on the component cut out by $k-1$ random hyperplanes. The factorization of the curve tells the decomposition of the solution component. Therefore, we restrict our explanation of using the linear trace to the case of a curve in the plane.

Suppose we have three points in the plane obtained as (projections of) witness points from some polynomial system. If the monodromy found loops between those points, then we know that these points lie on an irreducible factor of degree at least three. Whence our question: is this irreducible factor on which the given three points lie of degree three?

To answer this question we represent the factor by a cubic polynomial f in the form

$$\begin{aligned} f(x, y(x)) &= (y - y_1(x))(y - y_2(x))(y - y_3(x)) \\ &= y^3 - t_1(x)y^2 + t_2(x)y - t_3(x) \end{aligned}$$

Since $\deg(f) = 3$, $\deg(t_1) = 1$, so t_1 is the *linear trace*: $t_1(x) = c_1 x + c_0$.

We now proceed as follows. Via interpolation we find the coefficients c_0 and c_1. We first sample the cubic at $x = x_0$ and $x = x_1$. The samples are $\{(x_0, y_{00}), (x_0, y_{01}), (x_0, y_{02})\}$ and $\{(x_1, y_{10}), (x_1, y_{11}), (x_1, y_{12})\}$. To find c_0 and c_1 we then solve the linear system

$$\begin{cases} y_{00} + y_{01} + y_{02} = c_1 x_0 + c_0 \\ y_{10} + y_{11} + y_{12} = c_1 x_1 + c_0 \end{cases}$$

With t_1 we can predict the sum of the y's for a fixed choice of x. For example, samples at $x = x_2$ are $\{(x_2, y_{20}), (x_2, y_{21}), (x_2, y_{22})\}$, see Figure 8.19.
So our test consists in computing $t_1(x_2)$ in two ways:

$$c_1 x_2 + c_0 \stackrel{?}{=} y_{20} + y_{21} + y_{22}.$$

If the equality holds, then the answer to our question is yes.

Efficiency and numerical stability

The validation with the linear trace is fast. Therefore, our implementation does this validation each time a new loop with the monodromy algorithm

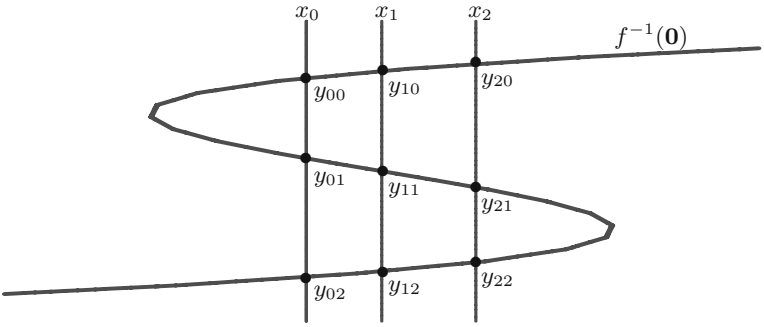

Fig. 8.19. The linear trace test on a planar cubic. To find the trace we interpolate through the samples at $x = x_0$ and $x = x_1$. Samples at $x = x_2$ are used in the test.

is found. Even as we do not know the locations of the singularities, practical experiences on many systems all lead to a rapid finding of permutations. While this approach is suitable for irreducible factors of very large degree (e.g., one thousand), strategies based purely on traces often perform better for smaller degrees.

Related to the efficiency is good numerical stability: if we can compute witness points with standard machine arithmetic, then we can also factor using standard machine arithmetic. This feature is very important when the accuracy of coefficients of the polynomial system is limited.

Exercise 8.3.2. Apply `phc -f` to factor

```
x**6 - x**5*y + 2*x**5*z - x**4*y**2 - x**4*y*z+x**3*y**3
  - 4*x**3*y**2*z + 3*x**3*y*z**2 - 2*x**3*z**3 + 3*x**2*y**3*z
  - 6*x**2*y**2*z**2 + 5*x**2*y*z**3 - x**2*z**4 + 3*x*y**3*z**2
  - 4*x*y**2*z**3 + 2*x*y*z**4+y**3*z**3 - y**2*z**4;
```

which is a polynomial in a format accepted by `phc`.

Exercise 8.3.3. Consider again the system of adjacent minors from Exercise 8.3.1. Determine the number of irreducible factors and their degrees.

See Chapter 9 for more on factorization methods.

8.4 Software and applications

8.4.1 Software for polynomial homotopy continuation

We agree with the statement: "It can be argued that the 'mission' of numerical analysis is to provide the scientific community with effective software tools." (taken from the preface to [GVL83]). Aside from our missionary intentions,

software has helped us in refining our algorithms, along the lines of the quote (from [Knu96]): "Another reason that programming is harder than the writing of books and research papers is that programming demands a significant higher standard of accuracy."

The software package PHCpack [Ver99a] is currently undergoing the transition from being a toolbox/black-box for various homotopy continuation methods to approximate all isolated solutions to a complete solving environment with capabilities to handle positive dimensional solution components efficiently, both in terms of computer operations and user manipulations. By the latter we hint at the search to find the right user interface, identifying the right data flow and trying to balance the toolbox with the black-box approach.

While PHCpack offered the first reliable implementation of polyhedral homotopies, its efficiency is currently surpassed by the implementations described in [GL00, GL03, LL01] and [DKK03, GKK+04, KK03b, TKF02]. To interact better with other codes, we are currently developing an interface from the Ada routines in PHCpack to routines written in C. Another (but related) interface concerns the interaction with computer algebra software. In [SVW03] we describe a very simple interface to Maple.

8.4.2 Applications

A benchmark suite for systems with positive dimensional solution components is gradually taking shape. Rather than listing summaries of a benchmark, we choose to treat two very typical applications: the cyclic n-roots problem from computer algebra and a special Stewart-Gough platform from mechanical design.

The cyclic n-roots problem. This problem is already interesting not only by its compact formulation and widespread fame in the computer algebra community, but also by known theoretical results concerning the number of isolated roots when n is prime [Haa96].

For $n = 8$, there are 16 one dimensional irreducible components: eight quadrics and eight curves of degree 16. While approximations to all 1,152 isolated cyclic 8-roots were found already in the first release of PHCpack, monodromy was needed to factor the curve of degree 144 into irreducibles. To compute all witness points for the cyclic 9-roots problem, the software of [LL01] was essential. While the factorization of a two dimensional component of degree 18 into six cubics posed no difficulty, the homotopy membership test was required to certify that among the 6,642 isolated ones 162 cyclic 9-roots occurred with multiplicity four. In addition, multi-precision arithmetic was used to confirm this result.

The isolated cyclic n-roots (up to $n = 13$, for which 2,704,156 paths were traced) can be found on the Internet[6] These roots have been computed with PHoM [GKK+04].

[6] `http://www.is.titech.ac.jp/~kojima/polynomials/cyclic13`.

A special Stewart-Gough platform. The Stewart-Gough platform is a parallel robot which attracted lots of interest from computational kinematicians and researchers in computer algebra. That the platform has forty isolated solutions was first established computationally by continuation [Rag93] and elimination methods [Laz92, Mou93], and later proved analytically [Hus96], [RV95], and [Wam96].

A six-legged platform (similar to the general Stewart-Gough platform) which permits motion was presented by Griffis and Duffy in [GD93] and first analyzed in [HK00]. It is called the Griffis-Duffy platform. Instead of forty isolated solutions we now consider a curve. In our formulation of the two cases we studied, twelve lines corresponded to degenerate cases deemed uninteresting from a mechanisms point of view. In the first case we were then left with one irreducible component of degree 28, while in the second case we found five components, four of degree six (one sextic was not reported in the analysis of [HK00]), and one component of degree four, see Figure 8.20.

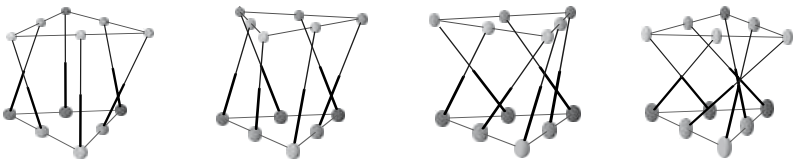

Fig. 8.20. One component of the Griffis-Duffy platform. Starting at the configuration at the left, we see the clockwise rotation of the end platform.

It is interesting to note that the running times for the factorization with the monodromy-traces method do not seem to depend on the particular geometry of the system, i.e.: the execution times are about the same in both cases, when we deal with one irreducible factor of high degree or with several factors of smaller degrees.

Acknowledgments

The authors thank Alicia Dickenstein and Ioannis Emiris for their invitation to present their work at the summer school. We are grateful to Dan Bates for his careful reading and comments. The revision benefited greatly from the stimulating questions from Olga Kashcheyeva, Anton Leykin, Yusong Wang, and Ailing Zhao at the MCS 595 graduate seminar. Some of the exercises were first presented at the RAAG summer school on Computer Tools for Real Algebraic Geometry, June 30-July 5, 2003, organized by Michel Coste, Laureano Gonzalez-Vega, Fabrice Rouillier, Marie-Françoise Roy, and Markus Schweighofer, whom we thank for their invitation.