

Unification and extension of intersection algorithms in numerical algebraic geometry

Jonathan D. Hauenstein* Charles W. Wampler †

July 27, 2016

Abstract

The solution set of a system of polynomial equations, called an algebraic set, can be decomposed into finitely many irreducible components. In numerical algebraic geometry, irreducible algebraic sets are represented by witness sets, whereas general algebraic sets allow a numerical irreducible decomposition comprising a collection of witness sets, one for each irreducible component. We denote the solution set of any system of polynomials $f : \mathbb{C}^N \rightarrow \mathbb{C}^n$ as $\mathcal{V}(f) \subset \mathbb{C}^N$. Given a witness set for some algebraic set $Z \subset \mathbb{C}^N$ and a system of polynomials $f : \mathbb{C}^N \rightarrow \mathbb{C}^n$, the algorithms of this article compute a numerical irreducible decomposition of the set $Z \cap \mathcal{V}(f)$. While extending the types of intersection problems that can be solved via numerical algebraic geometry, this approach is also a unification of two existing algorithms: the diagonal intersection algorithm and the homotopy membership test. The new approach includes as a special case the “extension problem” where one wishes to intersect an irreducible component A of $\mathcal{V}(g(x))$ with $\mathcal{V}(f(x, y))$, where f introduces new variables, y . For example, this problem arises in computing the singularities of A when the singularity conditions are expressed in terms of new variables associated to the tangent space of A . Several examples are included to demonstrate the effectiveness of our approach applied in a variety of scenarios.

Keywords. Numerical algebraic geometry, algebraic set, intersection, regeneration, witness set
AMS Subject Classification. 65H10, 68W30, 14Q99

1 Introduction

Numerical algebraic geometry concerns the solution of systems of polynomial equations using numerical methods, principally homotopy methods, also known as polynomial continuation. We take the ground field to be the complex numbers \mathbb{C} : the continuity and algebraic completeness of \mathbb{C} are essential for our algorithms. Given a polynomial system $F : \mathbb{C}^N \rightarrow \mathbb{C}^n$:

$$F(z_1, \dots, z_N) = \begin{bmatrix} F_1(z_1, \dots, z_N) \\ \vdots \\ F_n(z_1, \dots, z_N) \end{bmatrix}, \quad (1)$$

*Department of Applied and Computational Mathematics and Statistics, University of Notre Dame (hauenstein@nd.edu, www.nd.edu/~jhauenst). This author was partially supported by NSF grants DMS-1262428 and ACI-1460032, and Sloan Research Fellowship.

†General Motors Research and Development, Mail Code 480-106-RA2, 30500 Mound Road, Warren, MI 48090 (charles.w.wampler@gm.com, www.nd.edu/~cwampler1).

one wishes to describe its solution set, called the algebraic set defined by F , denoted as

$$\mathcal{V}(F) := \{z \in \mathbb{C}^N \mid F(z) = 0\}.$$

Early in the development of polynomial continuation, all attention was focused on finding just the isolated solutions of a given polynomial system. In these methods, a homotopy is constructed along with a set of starting solutions. A homotopy path emanates from each start point, and the set of all endpoints of these paths is guaranteed to include all the isolated solutions of the target polynomial system. See, for example, [33] for a treatise from the early era. Major advances in methods for constructing homotopies with fewer paths, which thereby generally reduce the computational cost of numerically tracking the paths, include multi-homogeneous formulations [34], linear-product homotopies [56], and polyhedral homotopies [25, 57]. The polyhedral methods make maximal use of whatever sparsity exists in the list of monomials appearing in the polynomials. While this gives the least number of homotopy paths among all the known single-shot homotopy methods, the homotopy construction phase depends on the computation of the mixed volume of the system [28, 31], potentially a computationally complex problem.

Numerical algebraic geometry was founded [50] on the recognition that often one may be interested in positive-dimensional solution sets (curves, surfaces, etc.). A methodology was developed to discover such sets by applying polynomial continuation to find witness points on the sets, these points having been isolated by intersecting the positive dimensional set with a generic linear space of complementary dimension. This idea led to the concept of a *witness set*, to be described precisely later, now considered the fundamental construct of the field. The computation of a numerical irreducible decomposition, first accomplished in [44], produces a witness set for each irreducible component of $\mathcal{V}(F)$. A concise review of the main concepts and vocabulary of the field, such as witness sets, irreducible components, and the numerical irreducible decomposition, are included here in §2; for more complete treatments, see [9, 51].

Software available for computing just isolated solutions includes Hom4PS [28], HomPack [36], and POLSYS_GLP [53]. Packages that compute both isolated solutions and witness sets for positive-dimensional solution sets are PHCpack [55] and Bertini [8]. The algorithms reported here have been tested using extensions to Bertini.

Soon after the crystallization of the witness set idea, it was seen that it can be advantageous to solve systems in stages. The diagonal intersection methods of [42, 43] compute the intersection of a pair of algebraic sets, each given by a witness set. Using this approach, one may introduce one-by-one the hypersurfaces defined by the individual polynomials in the system, maintaining a numerical irreducible decomposition of the intersection at each stage [48]. This kind of idea was given a new twist with the invention of regeneration methods [20, 21]. Both of these methods can be streamlined if one only wants the isolated solutions of the final system. As detailed in [23], to obtain a full irreducible decomposition, including positive-dimensional sets, the diagonal methods have to be completed using isosingular theory [22], based on Thom-Boardman singularities [3, 12].

The new algorithm described in this article generalizes the diagonal intersection problem and makes use of regeneration to solve it. The main problem we solve is to intersect a pure-dimensional algebraic set, $Z \subset \mathbb{C}^N$, given by a witness set, with $\mathcal{V}(F)$, for a polynomial system F as in (1). The result is a numerical irreducible decomposition of $Z \cap \mathcal{V}(F)$. Variants include the case when $Z = A \times B$ and we begin with witness sets for algebraic sets $A \subset \mathbb{C}^M$ and $B \subset \mathbb{C}^N$. In particular, given a witness set for $A \subset \mathcal{V}(G(x))$, this allows us to extend the solution to $(A \times \mathbb{C}^N) \cap \mathcal{V}(F(x, y))$, where $F : \mathbb{C}^M \times \mathbb{C}^N \rightarrow \mathbb{C}^n$ is a polynomial system that involves new variables, y , not present in $G(x)$. Since our approach extends from A using regeneration, we call it *regeneration extension*.

Our examples aim to show the utility of these new capabilities. At the same time, we note the unifying nature of this formulation, as both the existing diagonal homotopy method and the existing homotopy membership test can be viewed as special cases.

The enhanced ability to solve a variety of intersection problems gives great flexibility in how one may approach a problem and often has a large impact on the computational cost of doing so. It is quite common for systems arising in applications and in mathematical research to have solution sets composed of several irreducible components. Having computed a numerical irreducible decomposition for a system under study, the analyst often finds that some of the components are of further interest while it is desired to exclude the others from further investigation. This is because often some components are degenerate or nonphysical in some sense specific to the problem at hand. Our method allows one to compute intersections with just the interesting components while ignoring the rest. Our examples will illustrate how this can sometimes lead to considerable savings.

Once a witness set is available for an algebraic set, one may wish to investigate it further by finding points on it that satisfy extra conditions. One possibility is to find its singularities by intersecting the set with the conditions for a point on the set to have a tangent space of higher dimension than the set itself. With our new algorithm, one is free to write the singularity conditions in terms of new variables associated to the tangent directions rather than relying exclusively on determinants of matrices of partial derivatives in the original variables. The additional polynomials might instead define other “critical” conditions on the set, such as computing critical points of the distance to a fixed point, e.g., [4, 17, 40] (see also [14]), and critical points of a projection, e.g., [11, 32]. We will also illustrate our algorithm at work on this kind of problem.

Another application of the new algorithm arises in finding exceptional sets by taking sequences of fiber products, as in [52]. Not only does each new fiber product introduce a new set of variables, but also symmetry groups act on the irreducible components that arise. The ability to pick out one component for further investigation avoids the combinatorial explosion of fiber products that would otherwise ensue. This topic is too complicated to take up here, but it will be addressed in a forthcoming article. For now, we simply note that the kind of flexibility provided by the algorithm given here is crucial to tackling any but the simplest problems concerning exceptional sets.

After reviewing some background on the numerical irreducible decomposition and witness sets in §2, we state the main problems in §3 and begin the new developments with a simple algorithm in §4 for generating a witness set for $Z = A \times B$ given witness sets for A and B . Then, in §5, we show how to generate witness supersets for the pure-dimensional components of $Z \cap \mathcal{V}(F)$, followed by a description of how to complete the numerical irreducible decomposition of $Z \cap \mathcal{V}(F)$ in §6. In §7, we show in detail how the general machinery specializes to existing algorithms for diagonal intersection and for membership testing, as well as the specialization to the extension problem. Finally, in §8, we discuss some examples illustrating the efficacy of the new approach before recapping the contribution in §9.

2 Background

The following includes definitions of *dimension*, *degree*, *irreducible components*, *pure-dimensional components*, the *irreducible decomposition*, *witness sets*, and the term *generically reduced*. We also briefly discuss mathematical rigor and reliability in relation to numerical algebraic geometry. The reader familiar with such topics may wish to skip to the next section.

2.1 Definitions

An *algebraic set* is any set that can be written as $\mathcal{V}(F)$ for some polynomial system F , as in (1). For any algebraic set Z , we denote its dimension by $\dim Z$, defined as the maximum number, say k , such that k general hyperplanes meet Z .¹

An *irreducible algebraic set* is an algebraic set that cannot be expressed as the union of a finite number of its proper algebraic subsets. Every algebraic set can be expressed as a union of its irreducible components, that is, every algebraic set has an *irreducible decomposition* of the form

$$\mathcal{V}(F) = \bigcup_{i=0}^D Z_i, \quad Z_i = \bigcup_{j \in \mathcal{J}_i} Z_{ij}, \quad (2)$$

where each Z_{ij} is an irreducible component with $\dim Z_{ij} = i$, \mathcal{J}_i is a finite index set (possibly empty for $i < D$), $Z_{ij} \not\subset Z_{k\ell}$ for $(i, j) \neq (k, \ell)$, and $D = \dim \mathcal{V}(F) \leq N$. The irreducible decomposition is unique up to renumbering the components.

A *pure-dimensional* subset of $\mathcal{V}(F)$ is the union of any subset of irreducible components of $\mathcal{V}(F)$ having the same dimension. That is, using the notation of (2), a pure-dimensional subset of $\mathcal{V}(F)$ is any set of the form $\bigcup_{j \in \mathcal{I} \subset \mathcal{J}_i} Z_{ij}$ for some i and index subset \mathcal{I} . A *pure-dimensional component* of $\mathcal{V}(F)$ is simply any one of the Z_i , the union of all the i -dimensional irreducible components. Thus, we say that $\mathcal{V}(F) = \bigcup_{i=0}^D Z_i$ is the *pure-dimensional decomposition* of $\mathcal{V}(F)$.

For an irreducible algebraic set A , $\dim A$ general hyperplanes intersect A in a finite number of points, which number is defined as the *degree* of A , denoted $\deg A$. More precisely, the number of isolated points in $A \cap \mathcal{V}(L_A)$ is maximized on a dense, Zariski-open subset of affine linear mappings L_A from \mathbb{C}^N to $\mathbb{C}^{\dim A}$. Each L_A achieving this maximum number, denoted $\deg A$, is called *generic*. The degree of a pure-dimensional algebraic set is similarly defined and is equal to the sum of the degrees of its irreducible components.

The aforementioned definitions of dimension and degree lead naturally to the concept of a *witness set*, the basic construct of numerical algebraic geometry. For an irreducible algebraic set A , the witness set for A consists of three entries:

1. a polynomial system, f_A , called a *witness system*, such that A is an irreducible component of $\mathcal{V}(f_A)$;
2. a generic affine linear polynomial system, $L_A : \mathbb{C}^N \rightarrow \mathbb{C}^{\dim A}$; and
3. numerical approximations to the *witness point set* $W_A = A \cap \mathcal{V}(L_A)$, which is a finite set of $\deg A$ points.

Accordingly, we write a witness set for A as the ordered set $\mathcal{A} = \{f_A, L_A, W_A\}$. Clearly, the definition of a witness set can also be applied to pure-dimensional algebraic sets, in which case the witness point set is just the union of the witness point sets of its irreducible components. In this case, we can use the same affine linear polynomial system for all irreducible components.

A *numerical irreducible decomposition* of the solution set $\mathcal{V}(F)$ consists of one witness set for each of its irreducible components. The computation of such a decomposition proceeds in three stages, the first of which is the generation of a *witness point superset* for each pure-dimensional component. A witness point superset for Z_i is of the form $\widehat{W}_i = W_i \cup J_i$, where W_i is a witness

¹More precisely, for algebraic set $Z \subset \mathbb{C}^N$, $\dim Z$ is the maximum number k such that there exists a dense Zariski-open subset $U \subset \mathbb{C}^{k \times N} \times \mathbb{C}^k$ so that for every $\{A, b\} \in U$, $Z \cap \mathcal{V}(Az + b) \neq \emptyset$.

point set for Z_i and J_i consists of a finite number of points contained in $\bigcup_{\ell>i} Z_\ell$. The full process of computing a numerical irreducible decomposition for $\mathcal{V}(F)$ consists of the following three steps:

1. find a witness point superset \widehat{W}_i for each pure-dimensional component Z_i of $\mathcal{V}(F)$;
2. for each of these, eliminate the “junk set”, J_i , from \widehat{W}_i to obtain W_i ; and
3. partition W_i into witness point sets for the irreducible components of dimension i .

The impetus for defining and using pure-dimensional sets in this article is that it allows us to avoid in many cases the third step of decomposing pure-dimensional sets into their irreducible components.

Several methods apply to each of the three steps. Witness supersets can be computed using the cascade algorithm of [49], the dimension-by-dimension algorithm of [51, Chap. 13], or the regenerative cascade algorithm of [21]. The identification of “junk sets” J_i can be accomplished using a local dimension test [7] or, if the multiplicity depth is too large for that to be practical, by using the homotopy membership test [45] against all the higher-dimensional solution components. The break-up into irreducible components is done with monodromy [47] backed up by linear trace tests [46]. A general reference to these techniques is [51].

Consider an irreducible component A of $\mathcal{V}(f_A)$ with witness set as above. For any $a \in W_A$, we know that a is an isolated point in $\mathcal{V}(f_A, L_A)$. Let $J_A(a)$ be the Jacobian matrix of the system $\{f_A, L_A\}$ evaluated at a . If $J_A(a)$ has full column rank, i.e., $\dim \text{null } J_A(a) = 0$, then A is said to be *generically reduced* with respect to f_A , or equivalently, to have multiplicity one. Otherwise, A is said to be *generically nonreduced*, or equivalently, to have multiplicity greater than one. In this article, the main focus is on the generically reduced case. As discussed in §5.1, this does not represent a loss of generality, because through the use of *deflation* techniques, problems involving generically nonreduced algebraic sets can be solved using the generically reduced case.

2.2 On rigor and robustness

We note two caveats that generally apply to algorithms of numerical algebraic geometry, including this contribution.

First, the main algorithms all succeed with probability one in a theoretical sense, but in practice, this property is only approximated. That is to say, the algorithms involve one or more parameters living in an algebraic parameter space, and the good properties of the algorithm hold for any choice of parameters in a dense Zariski-open subset of that space, i.e., for all parameter choices that avoid some algebraic subset of the parameter space. An example of a good property is that all the homotopy paths are nonsingular and do not cross. In practice, we choose parameter values with a random number generator, which almost certainly avoids the bad set but does not absolutely guarantee it.

Second, the theoretical algorithms are implemented in floating point arithmetic, which introduces numerical error and can lead to algorithm failure or inaccurate results. To minimize these effects without an undue sacrifice in computational cost, the Bertini software package uses adaptive multiprecision, extending the number of floating point digits used according to the needs of the problem. These tactics do not guarantee success, but with some care in avoiding numerically ill-conditioned presentations of systems to be solved, one may expect high reliability. In some cases, results obtained in floating point can be verified in a mathematically rigorous sense by applying methods, such as alpha theory [5, Chap. 8] or interval arithmetic [41], that carefully bound all numerical errors. Without such rigorous, and often expensive, techniques, our algorithms must regularly make a decision of the form: given a numerical method of approximating some exact

value x^* by a floating point number \tilde{x} , is $f(x^*) = 0$? By evaluating \tilde{x} and $f(\tilde{x})$ in high precision, such decisions can be made with a high probability of correctness but shy of rigor.

The situation is analogous to the problem in numerical linear algebra of determining the rank of a matrix. One of the best tools for this is the singular value decomposition (SVD). When the SVD of a rank deficient matrix is computed in floating point, the rank deficient directions typically have associated singular values that are evaluated as very small numbers, near but not exactly zero. Consequently, it can be hard to distinguish between an ill-conditioned matrix and a truly singular one. If one increases the precision of the calculations, the numerical evaluations of singular values that are truly zero become smaller in direct response, while nonzero ones stabilize. In this manner, the determination of the rank of a matrix can be made quite robust but not rigorous, as there is no *a priori* lower bound on the size of a nonzero singular value. Implementations of algorithms in numerical algebraic geometry cannot avoid this issue, and adaptive multiprecision is the best response known at present.

3 Main Problems

We now state more precisely the main problems treated by our algorithm.

Problem 1 (Main Problem) *Given a witness set for a pure-dimensional algebraic set $Z \subset \mathbb{C}^N$ and a polynomial system $f : \mathbb{C}^N \rightarrow \mathbb{C}^n$, compute a numerical irreducible decomposition of $Z \cap \mathcal{V}(f)$.*

The diagonal intersection methodology mentioned in §1 can be used to solve this problem, albeit in a less direct fashion than we will soon describe. Reduction to the diagonal is based on the fact that for algebraic sets $A, B \subset \mathbb{C}^N$, their intersection, $A \cap B \subset \mathbb{C}^N$, is isomorphic to the set

$$\{(x, y) \in \mathbb{C}^{2N} \mid x \in A, y \in B, x = y\} \subset A \times B.$$

(The set $\mathcal{V}(x - y)$ is the *diagonal* of \mathbb{C}^{2N} , hence the name of the method.) The approach cascades down through all the possible dimensions where solution components of $A \cap B$ could exist, producing witness supersets for the pure-dimensional components of $(A \times B) \cap \mathcal{V}(x - y) \subset \mathbb{C}^{2N}$. After mapping these sets to \mathbb{C}^N using $(x, y) \mapsto (x)$, isosingular deflation methods are used in producing the witness systems for the irreducible components of $A \cap B$.

To apply the diagonal approach to our Main Problem, one must first compute the pure-dimensional decomposition of $\mathcal{V}(f) = \bigcup_{i=0}^{\dim \mathcal{V}(f)} V_i$. Then, by reduction to the diagonal, for each nonempty V_i , one computes witness supersets for $(Z \times V_i) \cap \mathcal{V}(x - y) \subset \mathbb{C}^{2N}$ for $x \in Z, y \in V_i$.

Our new approach simplifies the solving process by eliminating the decomposition of $\mathcal{V}(f)$ and proceeding directly with a computation in \mathbb{C}^N instead of doubling up to \mathbb{C}^{2N} . This is one of the main features of our new approach.

The answer to the Main Problem is the intersection of Z with the whole of $\mathcal{V}(f)$. In the case that one wishes to single out specific irreducible components of $\mathcal{V}(f)$, one must decompose $\mathcal{V}(f)$ to even ask the question. Then, one is forced into relying on reduction to the diagonal. Yet, this is but a special case of the following.

Problem 2 (Cartesian-Product Intersection) *Given witness sets for pure-dimensional algebraic sets $A \subset \mathbb{C}^M$ and $B \subset \mathbb{C}^N$ and a polynomial system $f : \mathbb{C}^{M+N} \rightarrow \mathbb{C}^n$, compute a numerical irreducible decomposition of $(A \times B) \cap \mathcal{V}(f)$.*

The reduction to the diagonal is a special case of this wherein $M = N$ and $f(x, y) = x - y$.

Obviously, the Cartesian-product intersection problem can be generalized, for example, to computing $(A \times B \times C) \cap \mathcal{V}(f)$ and beyond, and our methods easily handle this. The key step is to use the given witness sets to compute a single witness set for $Z = A \times B$ (or $Z = A \times B \times C$, etc.), thus reducing Cartesian-product intersection to the Main Problem.

Finally, we have the extension problem.

Problem 3 (Extension Problem) *Given a polynomial system $f : \mathbb{C}^{M+N} \rightarrow \mathbb{C}^n$ and a witness set for a pure-dimensional algebraic set $A \subset \mathbb{C}^M$, compute a numerical irreducible decomposition of $(A \times \mathbb{C}^N) \cap \mathcal{V}(f)$.*

This kind of problem arises naturally when A is part of the solution set of $g(x) = 0$, and one subsequently wishes to impose the conditions $f(x, y) = 0$ on A , where y is a set of new variables. Clearly, this is a special case of the Cartesian-product problem. We single it out because it arises frequently in applications and call our approach to solving this problem the *regeneration extension*.

Before proceeding to an algorithm for the Main Problem, we first show how to convert Cartesian-product intersections, and hence extension problems, into the Main Problem.

4 Cartesian products of irreducible sets

As noted above, it is useful to apply the capability to compute $Z \cap \mathcal{V}(f)$ to situations where $Z = A \times B$ for irreducible algebraic sets A and B . One could just as well consider the case where A and B are pure-dimensional but not necessarily irreducible, in which case Z is the union of the Cartesian products of all pairs of irreducible components taking one from A and one from B . With this, all of the same results discussed below in the irreducible case hold for pure-dimensional sets by applying the arguments to each pair. There is nothing in the procedures that follow that requires the witness point set for a pure-dimensional set to be decomposed according to irreducible components.

Let $A \subset \mathbb{C}^{N_A}$ and $B \subset \mathbb{C}^{N_B}$ be irreducible algebraic sets with $Z = A \times B$. Let x, y , and z be the coordinates on \mathbb{C}^{N_A} , \mathbb{C}^{N_B} , and $\mathbb{C}^{N_A+N_B}$, respectively, with $\bar{z} = \begin{bmatrix} z \\ 1 \end{bmatrix}$. We have the following facts:

1. Z is an irreducible algebraic set;
2. $\dim Z = \dim A + \dim B$;
3. $\deg Z = \deg A \cdot \deg B$; and
4. there exists a dense Zariski open subset U of $\mathbb{C}^{\dim Z \times (N_A+N_B+1)}$ such that for all $P \in U$ the set $W_Z := Z \cap \mathcal{V}(P \cdot \bar{z})$ consists of $\deg Z$ isolated points.

We know there exists polynomial systems $f_A(x)$ and $f_B(y)$ such that A and B are generically reduced with respect to f_A and f_B , respectively, with $A = \mathcal{V}(f_A)$ and $B = \mathcal{V}(f_B)$. For the system $f_Z(x, y) = \{f_A(x), f_B(y)\}$, one has $Z = \mathcal{V}(f_Z)$, e.g., see [16, Prop. 3.10] showing that Z is an algebraic set. Additionally, this observation immediately yields that $\dim Z = \dim A + \dim B$, e.g., see [27, Thm. 5.15]. For general $a \in A$ and $b \in B$, since the null space of the Jacobian matrix of f_Z evaluated at (a, b) , namely $\text{null } Jf_Z(a, b)$, also has dimension $\dim A + \dim B = \dim Z$, Z is generically reduced with respect to f_Z . It is clear that this observation yields $\deg(A \times \mathbb{C}^{N_B}) = \deg A$ and $\deg(\mathbb{C}^{N_A} \times B) = \deg B$.

The irreducibility of Z follows directly from the irreducibility of A and B as follows. Suppose that Z_1 and Z_2 are algebraic sets with $Z = Z_1 \cup Z_2$. For every $a \in A$, consider the irreducible set $B_a = \{a\} \times B$. Since $B_a = (B_a \cap Z_1) \cup (B_a \cap Z_2)$, the irreducibility of B_a implies that $B_a \subset Z_1$ or $B_a \subset Z_2$. Therefore, $A = A_1 \cup A_2$ where

$$A_i = \{a \in A \mid B_a \subset Z_i\}.$$

Since A is irreducible, it follows that $A_1 = A$ or $A_2 = A$ so that $Z_1 = Z$ or $Z_2 = Z$.

Suppose $\mathcal{L}_A \subset \mathbb{C}^{N_A}$ and $\mathcal{L}_B \subset \mathbb{C}^{N_B}$ are general linear spaces of codimension $\dim A$ and $\dim B$, respectively. Then, $\mathcal{L} = \mathcal{L}_A \times \mathcal{L}_B$ is a linear space of codimension $\dim A + \dim B = \dim Z$ and

$$Z \cap \mathcal{L} = (A \cap \mathcal{L}_A) \times (B \cap \mathcal{L}_B)$$

consists of $\deg A \cdot \deg B$ points. This shows $\deg Z \geq \deg A \cdot \deg B$. Bézout's Theorem provides the other inequality, namely

$$\deg(A \times B) = \deg((A \times \mathbb{C}^{N_B}) \cap (\mathbb{C}^{N_A} \times B)) \leq \deg(A \times \mathbb{C}^{N_B}) \cdot \deg(\mathbb{C}^{N_A} \times B) = \deg A \cdot \deg B.$$

The final follows from the previous and the Slicing Theorems [51, Thms. 13.2.1, 13.2.2].

In particular, this proof of the third fact shows that intersecting with a linear space constructed by the Cartesian product of general linear spaces in each set of variables produces the same number of points as with a general linear space in all variables. Hence, a witness set for $Z = A \times B$ can be easily constructed from witness sets for A and B as shown in the following lemma.

Lemma 4.1 *Let $\{f_A, L_A, W_A\}$ and $\{f_B, L_B, W_B\}$ be witness sets for irreducible and generically reduced algebraic sets $A \subset \mathbb{C}^{N_A}$ and $B \subset \mathbb{C}^{N_B}$, and let $Z = A \times B$. Also, let $\bar{z} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$, where x and y are coordinates on \mathbb{C}^{N_A} and \mathbb{C}^{N_B} , respectively. Then, for P chosen at random from $\mathbb{C}^{(\dim A + \dim B) \times (N_A + N_B + 1)}$, the homotopy*

$$h_0(x, y, t) = Z \cap \mathcal{V} \left(t \begin{bmatrix} L_A(x) \\ L_B(y) \end{bmatrix} + (1-t)(P \cdot \bar{z}) \right) \quad (3)$$

starting with the $\deg A \cdot \deg B$ points $W_A \times W_B$ at $t = 1$ is, with probability one, a complete homotopy (in the sense defined in [20]) for $Z \cap \mathcal{V}(P \cdot \bar{z})$. Moreover, letting W_Z be the set of endpoints of this homotopy at $t = 0$, $\{(f_A(x), f_B(y)), P \cdot \bar{z}, W_Z\}$ is a witness set for Z .

Proof. A complete homotopy for Y means that the paths are all trackable (exist, are continuous, and advance strictly monotonically with respect to $t \in (0, 1]$; see [20]) and their set of endpoints (the limits as $t \rightarrow 0$ of the paths) include all the isolated points in Y , where $Y = Z \cap \mathcal{V}(P \cdot \bar{z}) = W_Z$. From the facts above, we know that W_Z is a set of $\deg Z = \deg A \cdot \deg B$ points for generic $P \in \mathbb{C}^{(\dim A + \dim B) \times (N_A + N_B + 1)}$. By the parameter homotopy theorem [51, Thm. 7.1.1] (also [35]) and Lemma 7.1.2 of [51], homotopy h_0 defines $\deg Z$ trackable solution paths for $t \in [0, 1)$ starting at W_Z with endpoints that include all isolated roots of $H(x, y, 1) = 0$. But at $t = 1$, $h_0(x, y, 1) = 0$ has the $\deg Z$ isolated roots $W_A \times W_B$, so these must be the endpoints of the homotopy paths going from $t = 0$ to $t = 1$. As stated, the homotopy runs in the opposite direction, from $t = 1$ to $t = 0$, but the genericity of L_A and L_B along with the assumption that A and B are generically reduced imply that the points $W_A \times W_B$ are nonsingular roots of $h_0(x, y, 1) = 0$, hence the homotopy paths

are nonsingular for $t \in [0, 1]$. The final assertion that we have a witness set for Z is merely an observation that $\{(f_A(x), f_B(y)), P \cdot \bar{z}, W_Z\}$ fits the definition of a witness set. \square

In practice, one reduces to the “square” cases for numerical stability, namely polynomial systems $f_A : \mathbb{C}^{N_A} \rightarrow \mathbb{C}^{N_A - \dim A}$ and $f_B : \mathbb{C}^{N_B} \rightarrow \mathbb{C}^{N_B - \dim B}$ via randomization with Bertini’s Theorem. This is possible since A and B are only required to be irreducible components of their respective witness systems f_A and f_B . With this, h_0 from (3) is defined by the homotopy

$$h_0(x, y, t) = \left[\begin{array}{c} f_A(x) \\ f_B(y) \\ t \begin{bmatrix} L_A(x) \\ L_B(y) \end{bmatrix} + (1-t)(P \cdot \bar{z}) \end{array} \right], \quad (4)$$

which comprises $N_A + N_B$ polynomials in $N_A + N_B + 1$ variables (x, y, t) .

Remark 4.2 If the generically reduced assumption on A and B is dropped, the homotopy paths exist by the parameter homotopy theorem for isolated roots [51, Thm. 7.1.6]. After deflating the paths as in [20, §4], they can be tracked using a standard nonsingular path tracker. While this is a feasible approach, if either of A or B is generally nonreduced, we generally prefer to deflate before applying Lemma 4.1 (see §5.1).

Remark 4.3 For several irreducible algebraic sets, say A_1, \dots, A_k , it is obvious that Lemma 4.1 can be applied $(k-1)$ times in succession to obtain a witness set of $Z = A_1 \times \dots \times A_k$. This is not necessary since $\dim Z = \dim A_1 + \dots + \dim A_k$ and $\deg Z = \deg A_1 \cdots \deg A_k$ imply that a single homotopy of the form

$$h'_0(x_1, \dots, x_k, t) = Z \cap \mathcal{V}(t\{L_{A_1}(x_1), \dots, L_{A_k}(x_k)\} + (1-t)(P \cdot \bar{z})) \quad (5)$$

suffices, where the dimension of P is adjusted appropriately and the rest of the notation should be clear from context.

Remark 4.4 If one of the factors in the Cartesian product is Euclidean, say $B = \mathbb{C}^k$, then f_B is empty and W_B is the unique root of an arbitrary full-rank square linear system, $L_B : \mathbb{C}^k \rightarrow \mathbb{C}^k$. Everything proceeds in the same way as for a more general B .

5 Regenerative cascade

The objective of this section is to solve the following problem. This will serve as the first step towards solving Problem 1.

Problem 4 *Let $Z \subset \mathbb{C}^N$ be a pure-dimensional algebraic set represented by the witness set $\mathcal{Z} = \{f_Z, L_Z, W_Z\}$, and let $f : \mathbb{C}^N \rightarrow \mathbb{C}^n$ be a polynomial system. For each $i = 0, \dots, \dim Z$, compute a witness point superset, \widehat{W}_i , for the pure i -dimensional component of $Z \cap \mathcal{V}(f)$.*

We note that if $Z = \mathbb{C}^N$, this problem reduces to computing witness point supersets for the pure i -dimensional components $\mathcal{V}(f)$. The cascade algorithm of [49], the dimension-by-dimension algorithm of [51, Chap. 13], and the regenerative cascade algorithm of [21] were designed to solve this case. The difference here is to solve $f = 0$ on Z where Z is represented by the witness set \mathcal{Z} .

There are several ways to solve Problem 4 with existing algorithms. One option for the general case is to use one of the three algorithms mentioned in the previous paragraph to compute a numerical irreducible decomposition of $\mathcal{V}(f_Z, f)$. By using the membership test [45], one can then determine which ones are contained in Z . This may be inefficient particularly if $\mathcal{V}(f_Z, f)$ has many irreducible components not contained in Z .

Another option is to first compute numerical irreducible decompositions for both Z and $\mathcal{V}(f)$. Then, for each irreducible component of Z and $\mathcal{V}(f)$, use a diagonal intersection algorithm [42, 43] to compute a witness superset for each pure i -dimensional irreducible component of their intersection. The union of such witness supersets yields a solution to Problem 4. The approach of [42] performs this intersection using twice the number of variables while [43] employs intrinsic slicing to reduce the number of new variables used in the computation. In addition to the increased cost in performing the linear algebra routines for path tracking, this approach can have many redundant computations since each irreducible component of Z and $\mathcal{V}(f)$ are handled independently.

Our approach for solving Problem 4 is to modify the regenerative cascade algorithm of [21] to restrict all computations to set Z via the data in the witness set \mathcal{Z} . Like the existing algorithms, our generalization uses the following two concepts. For $x \in \mathbb{C}^N$, we say that x is a *nonsolution* of f if $f(x) \neq 0$. Additionally, the *rank* of f , denoted $\text{rank } f$, is defined as the rank of the Jacobian matrix of partial derivatives of f evaluated at a generic point of \mathbb{C}^N . We denote $L_Z = \{L_1, \dots, L_{\dim Z}\}$ where each L_i is a single linear polynomial defined on \mathbb{C}^N .

In brief, Algorithm 1 below proceeds by descending through the possible dimensions where pure dimensional components of $Z \cap \mathcal{V}(f)$ can exist, finding a witness point superset for each of these. The algorithm is written in terms of the co-dimension m of these components inside Z . The two main steps are the homotopy (6) that regenerates a start system of appropriate degree for stage m , and the homotopy (7) that solves stage m . At the end of each stage, the endpoints are sorted into solutions, which become the witness superset for that dimension, and nonsolutions, which are the start points for the next stage. Since the rank of f is a limit on the co-dimension of the intersection, the algorithm stops when it reaches either the zero dimension or exhausts the rank of f , whichever comes first.

Begin Algorithm 1

- Compute $r = \min(\text{rank } f, \dim Z)$.
- Choose a generic $r \times n$ matrix Λ and define $F = \{F_1, \dots, F_r\}$ as the polynomials formed as $F = \Lambda \cdot f$, where F and f are treated as column matrices. See Note 1 below for options.
- Set $m = 0$ and $U_0 = W_Z$.
- While true:
 1. Sort U_m into $U_m = \widehat{W}_{\dim Z - m} \cup X_m$, where $\widehat{W}_{\dim Z - m}$ are solutions and X_m are nonsolutions of f . See Note 2 below.
 2. If $m = r$, exit loop.
 3. Set $m = m + 1$. See Note 3 below.
 4. Form generic linear functions $L_{m,1}, \dots, L_{m,d_m}$, each a map from \mathbb{C}^N to \mathbb{C} , where $d_m = \deg F_m$. See Note 4 below.
 5. For $i = 1, \dots, d_m$:

– Track the solution paths of

$$h_{m,i}(z, t) = \{f_Z(z), F_1(z), \dots, F_{m-1}(z), \\ tL_m(z) + (1-t)L_{m,i}(z), \\ L_{m+1}(z), \dots, L_{\dim Z}(z)\} \quad (6)$$

from $t = 1$ towards $t = 0$, starting at the points X_{m-1} . Use a singular endgame, if necessary, to compute the endpoints (the limit as $t \rightarrow 0$) of each path. See Note 5 below.

– Let $T_{m,i}$ be the endpoints of the homotopy paths for $h_{m,i}$.

End for.

6. Let $T_m = \bigcup_{i=1}^{d_m} T_{m,i}$.

7. Track the solution paths of

$$h_m(z, t) = \{f_Z(z), F_1(z), \dots, F_{m-1}(z), \\ t \prod_{i=1}^{d_m} L_{m,i}(z) + (1-t)F_m(z), \\ L_{m+1}(z), \dots, L_{\dim Z}(z)\} \quad (7)$$

from $t = 1$ towards $t = 0$, starting at the points T_m . Use a singular endgame, if necessary, to compute the endpoints. See Note 5 below.

8. Let U_m be the finite endpoints of the homotopy paths for h_m .

End while.

- If $r < \dim Z$, then $\widehat{W}_i = \emptyset$ for all $i < \dim Z - r$.
- Return $\{\widehat{W}_0, \dots, \widehat{W}_{\dim Z}\}$.

End Algorithm 1

Notes:

1. The matrix Λ can be chosen with all entries below the main diagonal as zero. In addition, the entries in f can be reordered arbitrarily, and it is generally advantageous to place them in descending order by degree, in which case the degrees of the entries in F descend as well.
2. In practice, when using finite precision computation, care needs to be taken to distinguish between the solutions and nonsolutions. During each sorting call, one has the ability to refine the point to arbitrary accuracy, e.g., either through Newton's method or an endgame (discussed below in Note 4), to make this decision robust. Also see §2.2 for a general discussion.
3. As presented, Algorithm 1 proceeds equation-by-equation to compute witness supersets at each dimension. However, one may proceed subsystem-by-subsystem in order to skip dimensions not of interest, e.g., see §8.2.
4. Since $\{L_1, \dots, L_{\dim Z}\}$ are generic, it is acceptable to choose in every case $L_{i,1} = L_i$. Then, homotopy $h_{m,1}$ is trivial, and no computation is required to find $T_{m,1}$: it is exactly X_{m-1} .

5. Endgames, such as the Cauchy [38] and power series [37] endgame (see also [51, Chap. 10] and [26]) are used to compute the endpoint $z(0) = \lim_{t \rightarrow 0} z(t)$ of the homotopy path, $z(t)$, particularly if the endpoint is singular with respect to the corresponding homotopy at $t = 0$. Endgames work by building a local model of the path near $t = 0$ using data collected in the punctured neighborhood, i.e., away from $t = 0$. Being algebraic, the homotopy path always has a Puiseux (fractional power) series expansion convergent in the neighborhood of $t = 0$. The power series method estimates the coefficients of the Puiseux series to fit samples of $z(t)$ for t close to zero, with the constant term of the series being the estimate of $z(0)$. The Cauchy endgame instead computes a Cauchy integral of $z(t)$ for a closed path encircling the origin. (The winding number of the solution determines how many times the closed path in t must be tracked for the fiber above it, $z(t)$, to close up.) If one collects samples spaced uniformly around a circle centered on $t = 0$, the trapezoid integration rule gives $z(0)$ simply as the average of the points, an approach with excellent numerical properties, as discussed in [54].

We will say that Algorithm 1 solves Problem 4 if every homotopy path in the algorithm is trackable and the output is a valid set of witness point supersets for $Z \cap \mathcal{V}(f)$.

Theorem 5.1 *For random choices of all its generic coefficients, Algorithm 1 solves Problem 4 with probability one.*

Proof. The regenerative cascade method of [21] solves Problem 4 for the case that $Z = \mathbb{C}^N$ for some N . However, examination of the proof of that procedure with reference to the Simple Bertini Theorem for Systems [51, Thm. A.8.7] shows that it also holds if Z is any irreducible quasiprojective algebraic set. Furthermore, if Z is the union of irreducible quasiprojective sets of the same dimension, then there is no difference in carrying out the algorithm on one of these components at a time versus carrying it out on the whole set Z , that is, we do not need the witness point set W_Z to be decomposed into witness point sets for the irreducible components of Z . \square

5.1 Generically nonreduced sets

The statement of Problem 4 did not require Z to consist of the union of generically reduced irreducible sets $\mathcal{V}(f_Z)$, that is, Z might be a union of several irreducible solution components of $\mathcal{V}(f_Z)$ where some of these might have multiplicity larger than one. If each component of Z is generically reduced, then we have the desirable property that every nonsolution point in X_m and every point in T_m produced by the algorithm is nonsingular with respect to the corresponding polynomial systems. Then, the solution paths of every $h_{m,i}(z, t) = 0$ and $h_m(z, t) = 0$ are also nonsingular for $t \in (0, 1]$, so a nonsingular path tracker suffices.

In the case that Z has one or more irreducible components that is generically nonreduced, we can only say that the points X_m and T_m are isolated, and that the homotopy paths exist and are trackable in the sense defined in [20]. Some of the paths will be singular, that is, the Jacobian matrix of partial derivatives of the homotopy function with respect to z is not full-rank. As detailed in [20], trackable paths can be desingularized using a deflation operation, after which a nonsingular path tracker can be applied.

When faced with a generically nonreduced component of Z , an alternative is to deflate it first before applying the regenerative cascade algorithm. Methods for deflating a generically nonreduced irreducible set can be found in [22]. Since different irreducible components of Z may have different deflation sequences (equivalent to Thom-Boardman singularity sequences [3, 12] and summarized in §6), it is required to partially decompose Z based on identical deflation sequences. Irreducible

components with the same deflation sequence can be simultaneously deflated. The determinantal deflation procedure of [22], which does not introduce any auxiliary variables, is directly compatible with the regenerative cascade.

The other deflation procedures summarized in [22], e.g., building on the approach in [30], introduce new variables associated to tangent directions on the set. For these, the regenerative cascade requires minor modifications. If we suppose that Z is irreducible and generically nonreduced, then such a deflation produces a polynomial system g , an irreducible and generically reduced algebraic set $Y \subset \mathcal{V}(g)$, and a projection map π which is generically one-to-one on Y with $Z = \overline{\pi(Y)}$. Thus, one may attempt to perform computations on Y in place of Z . Following the approach of [18], one needs to change from considering paths which converge in Y to considering the paths for which the image under π converges. For example, for each $z \in Z \setminus \pi(Y)$, there is a path $\alpha : (0, 1] \rightarrow Y$ such that $z = \lim_{t \rightarrow 0} \pi(\alpha(t))$. However, $\alpha(t)$ must diverge as t approaches 0 since $z \notin \pi(Y)$. Here, only the endpoint of the image under π of the path defined by α was outside of $\pi(Y)$. By genericity, this is true for all paths arising in the regenerative cascade method as well.

5.2 Extrinsic and intrinsic homotopy

In both the homotopies $h_{m,i}$ of (6) and h_m of (7), the linear functions $L_{m+1}, \dots, L_{\dim Z}$ stand unperturbed as t varies. The same is true for any linear functions in the system F_Z or among the functions F_1, \dots, F_{m-1} at stage m of the cascade. Gathering all the unchanging linear functions into one linear subsystem of the homotopy, one may use linear algebra to compute the kernel of this subsystem before path tracking commences and then restrict computation in the path tracker to that subspace. Such an approach is said to be working intrinsically on the linear subspace. In contrast, an extrinsic method treats the linear functions just like any other polynomial in the system and, in essence, re-solves the linear part at each step of the path tracker. If there are enough linear functions present, then the intrinsic approach is more efficient, but if there are only a few unchanging linear functions, the extrinsic approach wins. When extrinsic wins, it is because the number of elements in the representation of a basis for the kernel is large, which raises the expense of working intrinsically. Intrinsic implementations of homotopies are discussed in [43, 20, 21] and specifics on assessing the trade-off between extrinsic and intrinsic formulations can be found in [23].

6 Completing the decomposition

The output of Algorithm 1 is a witness point superset, \widehat{W}_i , for each pure i -dimensional component of $Z \cap \mathcal{V}(f)$. As described in the introduction, two steps remain to produce a numerical irreducible decomposition, namely:

- eliminate the “junk sets,” J_i , from \widehat{W}_i to obtain W_i ; and
- partition W_i into witness point sets for the irreducible components of dimension i .

If Z is an irreducible component of $\mathcal{V}(f_Z)$ such that each irreducible component of $Z \cap \mathcal{V}(f)$ is an irreducible component of $\mathcal{V}(f_Z, f)$, then all of the standard numerical algebraic geometric methods apply. That is, the local dimension test [7] and the homotopy membership test [45] can be used to identify J_i , the irreducible components can be deflated at each witness point using positive-dimensional deflation techniques first described in [51, §13.3.2, §15.2.2], and monodromy [47] backed up by linear trace tests [46] can be used to decompose W_i . The following, motivated by [23, Ex. 2.0.2], highlights some shortcomings of these standard techniques in the context of intersection.

Example 6.1 Consider $f_Z(x, y, z) = (x + y + z)y$ with $Z = \mathcal{V}(x + y + z)$ and $f(x, y, z) = y$. Clearly, $Z \cap \mathcal{V}(f) = \mathcal{V}(x + z, y)$ is a line but $\mathcal{V}(f_Z, f) = \mathcal{V}(y)$ is a plane. In particular, for each $a \in \mathbb{C}$, the local dimension of $(a, 0, -a)$ with respect to $Z \cap \mathcal{V}(f)$ is 1, but is 2 with respect to $\mathcal{V}(f_Z, f)$. Thus, the local dimension test of [7] cannot be used for computing the “junk sets” J_i . Also, each $(a, 0, -a)$ is a smooth point on the plane $\mathcal{V}(f_Z, f) = \mathcal{V}(y)$ showing that standard deflation techniques applied to $\{f_Z, f\}$ at $(a, 0, -a)$ will fail to provide a polynomial system with the requisite irreducible component. Without such a system, the homotopy membership test, monodromy, and linear trace tests cannot be employed. Thus, we need to recast Problem 1, i.e., an intersection problem, in a way so that standard methods can be used. Example 6.3 demonstrates it for the setup above.

The key to overcoming these shortcomings and enabling the computation of a numerical irreducible decomposition for $Z \cap \mathcal{V}(f)$ is the theory of isosingular sets [22]. Let f_Z be the polynomial system given at the outset in the witness set for Z , i.e., Z is an irreducible component of $\mathcal{V}(f_Z)$. After completing the regeneration cascade, we have witness point supersets $\widehat{W}_i = W_i \cup J_i$. For each witness point $w \in W_i$ with corresponding irreducible component $X \subset Z \cap \mathcal{V}(f)$, we describe below how to use isosingular theory to construct a polynomial system f_w from w , f_Z , and f such that X is an irreducible component of $\mathcal{V}(f_w)$. At the largest k such that $\widehat{W}_k \neq \emptyset$, we know that $\widehat{W}_k = W_k$; that is, $J_k = \emptyset$. So we know that every point in W_k is a witness point, not junk, and therefore once f_w is in hand, monodromy, and linear trace tests discover which of the witness points form a witness set for X and likewise for all of the components at dimension k . Via the homotopy membership test, we can use these witness sets to eliminate from the lower-dimensional witness point supersets any junk points lying on sets of dimension k , and in particular, eliminate all junk from \widehat{W}_{k-1} . This leaves us in position to decompose W_{k-1} and to proceed in like fashion down through all the dimensions.

The key step of constructing system f_w depends on isosingular deflation, for which we provide only the key concepts here. We refer the interested reader to [22] for a general overview with [23] providing details related to diagonal intersection.

For a polynomial system G and a point $z \in \mathcal{V}(G) \subset \mathbb{C}^N$, $\text{Iso}_G(z)$ is an irreducible algebraic subset of $\mathcal{V}(G)$ containing z . Since the definition of $\text{Iso}_G(z)$ depends on the deflation sequence of G at z , we define this first and then construct $\text{Iso}_G(z)$. Let

$$\text{dnull}(G, z) := \dim \text{null } JG(z) = N - \text{rank } JG(z)$$

where $JG(z)$ is the Jacobian matrix of G evaluated at z and $\text{null } JG(z)$ is the right null space of $JG(z)$. Let $D(G, z)$ be the polynomial system consisting of G and the $(\text{rank } JG(z) + 1) \times (\text{rank } JG(z) + 1)$ minors of JG . Define $D^0(G, z) := G$ and, for $k \geq 1$, let $D^k(G, z)$ be the polynomial system obtained in this fashion after iterating k times. In particular, $D^k(G, z)$ is the polynomial system for the k^{th} isosingular deflation of G at z . The deflation sequence of G at z is

$$d_k(G, z) := \text{dnull}(D^k(G, z), z), \quad k = 0, 1, \dots$$

The closure (which is the same in both the Zariski topology and the standard Euclidean topology) of the set of points in $\mathcal{V}(G)$ that have the same deflation sequence as z with respect to G is an algebraic set that may decompose into several irreducible components. The irreducible set $\text{Iso}_G(z)$ is the unique such irreducible component that contains z . In particular, the deflation sequence of points in $\text{Iso}_G(z)$ is constant on a nonempty Zariski open subset. Such statement is true for any irreducible algebraic subset of $\mathcal{V}(G)$.

The following theorem, which is a slight generalization of [23, Thm. 5.1.1], describes deflating the irreducible components of $Z \cap \mathcal{V}(f)$.

Theorem 6.2 *Let $Z \subset \mathbb{C}^N$ consist of a union of irreducible components of $\mathcal{V}(f_Z)$. Suppose that $f : \mathbb{C}^N \rightarrow \mathbb{C}^n$ is a polynomial system, $F(x, y) = \{f_Z(x), f(y)\}$, $\Delta = \{(x, x) \mid x \in \mathbb{C}^N\}$, and $\pi(x, y) = x$. If $A \subset Z \cap \mathcal{V}(f)$ is an irreducible component, there exists a nonempty Zariski open set $U \subset A$ such that for all $a \in U$, A is an irreducible component of $\pi(\text{Iso}_F((a, a)) \cap \Delta)$. Thus, for a linear space $\mathcal{L} \subset \mathbb{C}^N$ of dimension $N - \dim A$ chosen randomly, then, with probability one, the witness point set $A \cap \mathcal{L}$ is contained in U .*

Proof. Since $\mathcal{A} = (A \times A) \cap \Delta$ is an irreducible algebraic set contained in $\mathcal{V}(F)$, there is a nonempty Zariski open set $\mathcal{U} \subset \mathcal{A}$ such that the deflation sequence with respect to F is constant on \mathcal{U} . Thus, $U = \pi(\mathcal{U})$ is a nonempty Zariski open subset of A with $\mathcal{U} = (U \times U) \cap \Delta$. Fix $a \in U \subset A$. Since $\mathcal{A} \subset \text{Iso}_F((a, a)) \cap \Delta$, there is an irreducible component $\mathcal{W} \subset \text{Iso}_F((a, a)) \cap \Delta$ with $W = \pi(\mathcal{W})$ such that $\mathcal{A} \subset \mathcal{W}$ and $A = \pi(\mathcal{A}) \subset \pi(\mathcal{W}) = W$. Thus, W is an irreducible component of $\pi(\text{Iso}_F((a, a)) \cap \Delta)$ containing A . The first part of the theorem will follow by simply showing that $A = W$.

Since $A \subset Z$ and Z is a union of irreducible components of $\mathcal{V}(f_Z)$, there is an irreducible algebraic set Z' which an irreducible component of $\mathcal{V}(f_Z)$ such that $A \subset Z' \subset Z$. Similarly, since $A \subset \mathcal{V}(f)$, there is an irreducible component $V' \subset \mathcal{V}(f)$ such that $A \subset V'$. Hence, $Z' \times V'$ is an irreducible component of $\mathcal{V}(F)$ with $(a, a) \in Z' \times V'$, i.e., $\text{Iso}_F((a, a)) \subset Z' \times V' \subset Z \times \mathcal{V}(f)$. Hence,

$$A \subset W \subset \pi(\text{Iso}_F((a, a)) \cap \Delta) \subset \pi((Z \times \mathcal{V}(f)) \cap \Delta) = Z \cap \mathcal{V}(f).$$

Since A is an irreducible component of $Z \cap \mathcal{V}(f)$, it follows that $A = W$.

If $\dim A = i$, the last statement follows from the the fact that, for a general linear subspace \mathcal{L}_i of codimension i , $A \cap \mathcal{L}_i = U \cap \mathcal{L}_i$. \square

Since each isosingular set is deflatable, Theorem 6.2 provides a prescription for constructing polynomial systems that can be used to complete the last two steps of computing a numerical irreducible decomposition of $Z \cap \mathcal{V}(f)$. Although we have no example, we cannot rule out the possibility that the corresponding irreducible component might be generically nonreduced with respect to the polynomial system constructed in this fashion. In such a case, one may simply use finitely many, bounded by the depth [13] and strictly bounded by the multiplicity [30], more steps of isosingular deflation or another deflation procedure for irreducible components, e.g., [51, §13.3.2, §15.2.2], to simplify to the generically reduced case.

The following example demonstrates using Theorem 6.2 on Example 6.1.

Example 6.3 *For f_Z , f , and Z as described in Example 6.1, $A = Z \cap \mathcal{V}(f) = \mathcal{V}(x + z, y)$ is irreducible. Theorem 6.2 allows us to derive this result algorithmically as follows. Suppose we have the witness point $w = (a, 0, -a) \in A$ and the polynomials f_Z and f . We wish to find a polynomial system f_w such that A is an irreducible component of $\mathcal{V}(f_w)$. Following the setup in Theorem 6.2, we first form*

$$F(x, y, z, x', y', z') = \begin{bmatrix} f_Z(x, y, z) \\ f(x', y', z') \end{bmatrix} = \begin{bmatrix} (x + y + z)y \\ y' \end{bmatrix}.$$

For general $a \in \mathbb{C}$, one can verify that the deflation sequence of $(a, 0, -a, a, 0, -a)$ with respect to F is $5, 3, 3, \dots$ with the corresponding 3-dimensional set

$$\text{Iso}_F((a, 0, -a, a, 0, -a)) = \{(b, 0, -b, \alpha, 0, \beta) \mid b, \alpha, \beta \in \mathbb{C}\}.$$

The polynomial system defining this isosingular set is formed by adding the 2×2 minors of JF to F . That is, $\text{Iso}_F((a, 0, -a, a, 0, -a))$ is an irreducible component of $\mathcal{V}(G)$ where

$$G(x, y, z, x', y', z') = \begin{bmatrix} (x + y + z)y \\ y' \\ y \\ x + 2y + z \\ y \end{bmatrix}.$$

Theorem 6.2 implies that A is an irreducible component of $\mathcal{V}(G(x, y, z, x, y, z))$; that is, $f_w(x, y, z) = G(x, y, z, x, y, z)$ suffices to form a witness set for A . It is easily seen that this can be simplified to just $f_w(x, y, z) = \{x + z, y\}$.

7 Special cases

We now return to elucidate some details of two special cases highlighted in the Introduction.

7.1 Diagonal intersection

To compute $A \cap B \subset \mathbb{C}^N$ for irreducible algebraic sets A and B described via witness sets, one may first apply Lemma 4.1 to obtain a witness set for $A \times B$ and then find $(A \times B) \cap \mathcal{V}(x - y)$ using Algorithm 1. There are special items to note about this case.

First, the system $f = x - y$ is linear and so is $F = \Lambda \cdot f$ in Algorithm 1. This has two implications:

- the homotopies $h_{m,i}$ of (6) are always trivial if we choose the option of Note 4 of Algorithm 1; and
- the intrinsic approach discussed in §5.2 is very effective for solving h_m in (7). In fact, only one linear function is perturbed in each h_m .

Second, since $(A \times B) \cap \mathcal{V}(x - y)$ and $A \cap B$ are isomorphic, their pure-dimensional components exist at the same dimensions. The largest possible dimension of an irreducible component of $A \cap B$ is $D = \min(\dim A, \dim B)$. Since Algorithm 1 starts at $\dim Z = \dim A + \dim B$ and descends, this means that the witness point supersets for dimension $D + 1, D + 2, \dots, \dim Z$ generated by the algorithm are empty. Instead of starting with Lemma 4.1 and proceeding through several unproductive stages of Algorithm 1, one can target directly the dimension D to avoid wasted computation. Let $d = \dim A + \dim B - D = \max(\dim A, \dim B)$. A suitable homotopy for targeting dimension D is

$$h_0''(x, y, t) = \begin{bmatrix} f_A(x) \\ f_B(y) \\ \gamma t \begin{bmatrix} L_A(x) \\ L_B(y) \end{bmatrix} + (1 - t) \begin{bmatrix} F_1(x, y) \\ \vdots \\ F_d(x, y) \\ L_{d+1}(x, y) \\ \vdots \\ L_{\dim A + \dim B}(x, y) \end{bmatrix} \end{bmatrix}, \quad (8)$$

where γ is chosen at random from the unit circle in \mathbb{C} . This homotopy starts from the points $W_A \times W_B$. In effect, this combines Lemma 4.1 with d levels of Algorithm 1—each of which would

have $\deg A \cdot \deg B$ homotopy paths—into just one homotopy with that number of paths. Since the target functions $F_1, \dots, L_{\dim A + \dim B}$ are all linear and since the linear system $\{L_A, L_B\}$ cuts out $\deg Z$ isolated points on Z at $t = 0$, the homotopy is seen to succeed with probability one using the Gamma Trick [51, Lemma 7.1.3]. (The gamma trick is a generalization of a technique first introduced in [34].)

Using these specializations, our algorithm for the Main Problem becomes an efficient approach to diagonal intersections. The intrinsic approach given here is simpler to implement and just as efficient than the original intrinsic formulation of [43].

7.2 Membership testing

Another special case is testing whether a given point, say x^* , is a member of an irreducible set A . Membership is true if $A \cap \{x^*\} = \{x^*\}$, else $A \cap \{x^*\} = \emptyset$. So this is a special case of the diagonal intersection: we wish to compute $(A \times \{x^*\}) \cap \mathcal{V}(x - x^*)$. Since $\dim\{x^*\} = 0$, one sees that (8) directly targets a system $\{F_1, \dots, F_{\dim A}\}$. But this system is just $\Lambda \cdot (x - x^*)$ for generic $\Lambda \in \mathbb{C}^{\dim A \times N}$, so the homotopy (8) becomes

$$h_0''(x, t) = \left[\begin{array}{c} f_A(x) \\ tL_A(x) + (1-t)\Lambda \cdot (x - x^*) \end{array} \right], \quad (9)$$

where the gamma trick is no longer needed since $L_A(x)$ is generic in the only variables appearing, namely x . Note that this one homotopy completes the entire procedure, as there are no more dimensions left for the cascade.

One can see that this test is exactly the homotopy membership test introduced in [47] and discussed in [51, §15.4]. While the approach of this paper does not improve on the existing membership test, it is unifying to see the test arise naturally as a special case of our treatment of diagonal intersection.

8 Examples

The following examples demonstrate the techniques on a variety of problems. For the timings reported, the examples used Bertini v1.5 [8] running either in serial on one of the sixteen cores of an AMD Opteron 6376 2.3 GHZ processor with 128 GB of memory or in parallel using four such processors for a total of 64 cores. All files needed to recreate the solving of these systems are available at www.nd.edu/~jhauenst/extension.

8.1 Comparison of the methods

In order to compare various methods for performing the first step of computing a numerical irreducible decomposition, i.e., computing witness point supersets for the pure-dimensional components, consider the first-order deflated variety [29] for the cyclic-4 system, which was considered in [29, Ex 5.5]. For $x, \lambda \in \mathbb{C}^4$, this variety is defined by

$$F(x, \lambda) = \left[\begin{array}{c} f(x) \\ Jf(x) \cdot \lambda \end{array} \right] \quad \text{where} \quad f(x) = \left[\begin{array}{c} x_1 + x_2 + x_3 + x_4 \\ x_1x_2 + x_2x_3 + x_3x_4 + x_4x_1 \\ x_1x_2x_3 + x_2x_3x_4 + x_3x_4x_1 + x_4x_1x_2 \\ x_1x_2x_3x_4 - 1 \end{array} \right].$$

dim	deg	# comp
1	2	2

dim	deg	# comp
4	1	9*
	6	1
	8	1

dim	deg	# comp
2	1	8
	4	2

Figure 1: Summary of the irreducible components related to the cyclic-4 system. All are generically reduced except (*) one linear space with respect to $Jf(x) \cdot \lambda$.

The irreducible components for $\mathcal{V}(f)$, $\mathcal{V}(Jf(x) \cdot \lambda)$, and $\mathcal{V}(F)$ are summarized in Figure 1.

Under the projection $(x, \lambda) \mapsto x$, eight of the ten irreducible components of $\mathcal{V}(F)$ project onto the eight embedded points defined by f . The other two, say V_1 and V_2 , project onto the two irreducible curves of $\mathcal{V}(f)$.

In this illustrative example, we use serial processing. First, we compute witness point supersets directly for $\mathcal{V}(F)$ using

- the dimension-by-dimension algorithm of [51, Chap. 13]: 212 seconds,
- the cascade algorithm of [49]: 146 seconds, and
- the regenerative cascade algorithm of [21]: 74 seconds.

Next, to use a diagonal intersection approach [42, 43], we first compute a numerical irreducible decomposition of $\mathcal{V}(f)$ and $\mathcal{V}(Jf(x) \cdot \lambda)$ using the regenerative cascade algorithm since $\mathcal{V}(F) = (\mathcal{V}(f) \times \mathbb{C}^4) \cap \mathcal{V}(Jf(x) \cdot \lambda)$. Each of these computations took under a second. As summarized in Figure 1, $\mathcal{V}(f)$ consists of two irreducible curves while $\mathcal{V}(Jf(x) \cdot \lambda)$ consists of eleven irreducible components of dimension 4. Since one of the linear spaces is generically nonreduced with respect to $Jf(x) \cdot \lambda$, we added a 4×4 minor of the Jacobian matrix of the system. In our test, the total time to perform the initial decompositions together with the $2 \cdot 11 = 22$ intersections using

- the diagonal intersection approach [42, 43]: 19 seconds.

The diagonal approach was faster than the direct methods by first decomposing the smaller systems $\mathcal{V}(f)$ and $\mathcal{V}(Jf(x) \cdot \lambda)$. However, in this case, one of the components of $\mathcal{V}(Jf(x) \cdot \lambda)$ needed to be deflated and some additional analysis was required, e.g., both V_1 and V_2 arise from two different intersections.

By using a regeneration extension approach via Algorithm 1, we first compute, in under a second, a decomposition of $\mathcal{V}(f)$ using the regenerative cascade algorithm and then extend to compute a decomposition of the first-order deflated variety $\mathcal{V}(F) = (\mathcal{V}(f) \times \mathbb{C}^4) \cap \mathcal{V}(Jf(x) \cdot \lambda)$. In our test, the total time using

- the regeneration extension approach: 4 seconds.

8.2 Four-bar linkages

The path synthesis problem for four-bar linkages is to design four-bar linkages whose coupler curve passes through the given points. In 1923, Alt [1] observed that the path synthesis problem described by nine points would, generically, yield finitely many four-bar linkages. Using homotopy

k	dim	deg
2	7	7
3	6	43
4	5	234
5	4	1108
6	3	3832
7	2	8716
8	1	10858
9	0	8652

Table 1: Summary of the nondegenerate components for designing four-bar linkages whose coupler curve passes through k specified points.

With this setup, $\mathcal{V}(F) \subset \mathbb{C}^{24}$ consists of $8652 = 2 \cdot 3 \cdot 1442$ isolated points due to the two-fold symmetry of swapping $\{x, \hat{x}, a, \hat{a}\}$ and $\{y, \hat{y}, b, \hat{b}\}$, and Roberts cognates. The other irreducible components of $\mathcal{V}(F)$ are all contained in one of the following 12 hyperplanes:

$$\begin{aligned} &\mathcal{V}(a - x), \quad \mathcal{V}(a - b), \quad \mathcal{V}(x - y), \quad \mathcal{V}(b - y), \quad \mathcal{V}(x), \quad \mathcal{V}(y), \\ &\mathcal{V}(\hat{a} - \hat{x}), \quad \mathcal{V}(\hat{a} - \hat{b}), \quad \mathcal{V}(\hat{x} - \hat{y}), \quad \mathcal{V}(\hat{b} - \hat{y}), \quad \mathcal{V}(\hat{x}), \quad \mathcal{V}(\hat{y}). \end{aligned} \quad (11)$$

Components lying in any of these hyperplanes have no physical meaning.

With parallel processing, we compute the 8652 points in $\mathcal{V}(F)$ using a diagonal approach [42, 43], as well as a regeneration extension approach that naturally proceeds subsystem-by-subsystem. That is, in both cases, we first compute $\mathcal{V}(F_1) \subset \mathbb{C}^{10}$ using the regenerative cascade algorithm, thereby solving the 2 point path synthesis problem. Then, we perform 7 subsequent intersections which solve the k point path synthesis problems for $k = 3, \dots, 9$. After each intersection, the points contained in one of the degenerate hyperplanes listed in (11) are identified and removed.

Since the first eight variables listed in (10) describe the four-bar linkages, we only want to compute the degrees with respect to the projection onto the first eight variables. This is accomplished by using general slices in only the first eight variables [19], i.e., for each $k = 2, \dots, 9$, we want to compute the number of isolated solutions of

$$F_1 = \dots = F_{k-1} = L_k = \dots = L_8 = 0 \quad (12)$$

not contained in a degenerate hyperplane (11) where each $L_j(x, \hat{x}, a, \hat{a}, y, \hat{y}, b, \hat{b})$ is a general linear polynomial. The results are summarized in Table 1.

With this setup, the 2-homogeneous Bézout count for $k = 2$ using the variable groupings $\{x, \hat{x}, a, \hat{a}, y, \hat{y}, b, \hat{b}\}$ and $\{\gamma_1, \hat{\gamma}_1\}$ is $2 \cdot (2 + 2) = 8$. Moreover, to use the regeneration extension approach to move from k to $k + 1$, following Note 4, each solution produces 3 paths to track.

The upper bound of 8 is not sharp since it is classically known that the $k = 2$ case has degree 7. Thus, by using a diagonal intersection approach, each nondegenerate solution at k produces 7 paths to be tracked at $k + 1$ compared with the regenerative cascade approach producing $3 + 8 = 11$ paths. This reduction in the number of paths is offset by an increased cost of performing the linear algebra computations since the diagonal approach adds additional variables. In our experiment, the total computational time was

- diagonal intersection: 206 minutes; and
- regeneration extension: 68 minutes.

8.3 Sphere packings

A sphere packing is an assemblage of rigid spheres in tangential contact such that no internal motion of the assembly is possible without breaking at least one contact. (Rigid motion of the whole assembly is allowed.) These are of interest, for example, in designing colloids formed by suspensions of microspheres. By coating the spheres with DNA, one can promote adhesion between them in patterns that bias the formation of certain packings, thereby controlling the properties of the resulting colloid [24]. One approach to predicting which clusters will form starts by determining all possible sphere packings of up to N spheres, with $N = 10$ being the largest case completed so far [2]. It is assumed that all spheres have the same diameter, which we take to be unit length.

A packing can be specified by numbering the spheres 1 to N and listing which ones are in contact. One representation of the contact list is an adjacency matrix, an $N \times N$ matrix with a 1 in the (i, j) -th element if spheres i and j touch, and 0 elsewhere. Renumbering the spheres while maintaining the same contacts produces an isomorphism between adjacency matrices. Finding all packings of N spheres requires sifting through all possible isomorphic groups to find those which are realizable (the solution is real and no two spheres intersect in more than a single point) and rigid.

Rigidity requires at least $3N - 6$ contacts, since each sphere center has three coordinates, each contact imposes one algebraic constraint, and the rigid assembly retains six degrees of freedom of rigid motion. Some packings have more than $3N - 6$ contacts; the smallest N where this occurs is $N = 10$ [2]. Also, it is possible for a packing to be singular in the sense that the assembly is a solution of multiplicity greater than one to the system of contact constraint equations. The smallest such packing occurs at $N = 9$ [2]. Arkus et al. call this packing “non-rigid,” which is reasonable considering that the singularity will allow a physical system, which can violate the contact conditions slightly, to vibrate in the direction associated to the null-space of the system Jacobian matrix. Thus, by their usage of the term, rigidity requires multiplicity one, while packings are assemblages that are isolated solutions of the system of contact equations.

As we shall see momentarily, it is common in these problems to find solution sets with components on which certain pairs of spheres coincide. These components are not realizable. In what follows, we call them degenerate.

Let us consider a single $N = 12$ case defined by the following contact pairs:

$$\{1-2, 1-3, 1-4, 1-5, 2-3, 2-4, 2-6, 3-5, 3-6, 3-7, 4-5, 4-6, 4-8, 5-7, 5-8, 5-10, 6-7, \\ 6-8, 6-12, 7-10, 7-11, 7-12, 8-9, 8-10, 8-12, 9-10, 9-11, 9-12, 10-11, 11-12\},$$

which is a graph with 12 vertices and 30 edges. With $3N - 6 = 30$ contacts, this case has the possibility of producing minimally constrained rigid packings.

Let v_i be the 3-vector of coordinates for vertex i . Since vertices 1, 2, and 3 contact pairwise, we set these to be the vertices of a unit equilateral triangle, i.e., $v_1 = (0, 0, 0)$, $v_2 = (1, 0, 0)$, and $v_3 = (1/2, \sqrt{3}/2, 0)$. We aim to compute v_4, \dots, v_{12} , a total of 27 variables, satisfying the remaining 27 contact relationships having the form

$$D_{i,j} := (v_i - v_j)^T(v_i - v_j) - 1 = 0$$

for contact pair i - j . Since v_1, v_2 , and v_3 are known, it is advantageous to rewrite as

$$D_{2,4} - D_{1,4} = D_{3,5} - D_{1,5} = D_{3,6} - D_{2,6} = \\ D_{1,4} = D_{1,5} = D_{2,6} = D_{3,7} = D_{4,5} = D_{4,6} = D_{4,8} = D_{5,7} = D_{5,8} = D_{5,10} = D_{6,7} = D_{6,8} = D_{6,12} = \\ D_{7,10} = D_{7,11} = D_{7,12} = D_{8,9} = D_{8,10} = D_{8,12} = D_{9,10} = D_{9,11} = D_{9,12} = D_{10,11} = D_{11,12} = 0,$$

yielding a system of 3 linear and 24 quadratic equations. Even though each of these 27 polynomials has at most 10 monomials, both the total degree and mixed volume for this system are equal to $2^{24} = 16,777,216$. Thus, the commonly used single-shot homotopies will require tracking 2^{24} paths.

Rather than relying upon a single-shot homotopy, our approach easily facilitates a flexible solving approach. The following shows one way to solve this $N = 12$ packing by looking at various subsystems. This procedure not only permits the solving of this particular $N = 12$ packing, it provides information, in the form of witness sets, for solving other related packing problems.

We start by breaking this $N = 12$ packing into two pieces by cutting along the closed path 5-8-6-7-5. This yields subgraph 1 with vertices $1, \dots, 8$ and edges

$$\{1-2, 1-3, 1-4, 1-5, 2-3, 2-4, 2-6, 3-5, 3-6, 3-7, 4-5, 4-6, 4-8, 5-7, 5-8, 6-7, 6-8\},$$

and subgraph 2 with vertices $5, \dots, 12$ and edges

$$\{5-7, 5-8, 5-10, 6-7, 6-8, 6-12, 7-10, 7-11, 7-12, 8-9, 8-10, 8-12, 9-10, 9-11, 9-12, 10-11, 11-12\}.$$

Since each has $N = 8$ spheres with $3N - 7 = 17$ edges, the corresponding solution sets must be at least one-dimensional. In fact, by relabeling subgraph 2 via

$$5 \rightarrow 7', \quad 6 \rightarrow 8', \quad 7 \rightarrow 5', \quad 8 \rightarrow 6', \quad 9 \rightarrow 2', \quad 10 \rightarrow 3', \quad 11 \rightarrow 1', \quad 12 \rightarrow 4',$$

one sees that these graphs are isomorphic. Let C_8 denote the set of all nondegenerate packings in 8 spheres satisfying the contact pairs described by subgraph 1. One can observe that each solution to our $N = 12$ packing problem is contained in

$$(C_8 \times C'_8) \cap \mathcal{V}(D_{5,6} - D_{7',8'}, D_{7,8} - D_{5',6'}) \tag{13}$$

where C'_8 is isomorphic to C_8 but has vertices $1', \dots, 8'$. The two additional equations ensure that the diagonals of the corresponding loops 5-8-6-7-5 and $7'-6'-8'-5'-7'$ are equal.

To compute C_8 , we could directly compute a numerical irreducible decomposition of the resulting polynomial system. Following a similar approach to the $N = 12$ packing, this results in a system of 3 linear and 11 quadratic polynomials in 15 variables. For example, based on using the regenerative cascade algorithm [21], it took 34 seconds using parallel processing to produce the following irreducible components:

- four two-dimensional components of degree 4 with $v_5 = v_6$;
- two-dimensional components of degree 32 and 48 with $v_3 = v_4$;
- four one-dimensional components of degree 2 with $v_1 = v_6$;
- four one-dimensional components of degree 4 with $v_1 = v_6$;
- four one-dimensional components of degree 4 with $v_2 = v_5$;
- nondegenerate one-dimensional component, call it A_1 , of degree 12;
- nondegenerate one-dimensional component, call it A_2 , of degree 14.

With this setup, $C_8 = A_1 \cup A_2$, which has degree 26.

Before considering another approach for computing C_8 , let's use this decomposition to solve the $N = 12$ packing problem. Due to the symmetry in (13), we need to compute

$$(A_i \times A_j) \cap \mathcal{V}(D_{5,6} - D_{7',8'}, D_{7,8} - D_{5',6'}) \text{ for } (i, j) = (1, 1), (1, 2), (2, 2).$$

Using Algorithm 1, we compute the following:

- $|(A_1 \times A_1) \cap \mathcal{V}(D_{5,6} - D_{7',8'}, D_{7,8} - D_{5',6'})| = 208$ in 34 seconds;
- $|(A_1 \times A_2) \cap \mathcal{V}(D_{5,6} - D_{7',8'}, D_{7,8} - D_{5',6'})| = 80$ in 30 seconds;
- $|(A_2 \times A_2) \cap \mathcal{V}(D_{5,6} - D_{7',8'}, D_{7,8} - D_{5',6'})| = 80$ in 45 seconds.

Of these 368 points, 144 are real with three additional checks needed to ensure a proper packing. First, testing that the spheres must only intersect in a point eliminates 140 of the 144 real solutions. These 4 remaining solutions all arise from $A_1 \times A_1$. Second, a congruence check is necessary since having all six distance pairs equal to one for vertices $(5, 6, 7, 8)$ and for $(5', 6', 7', 8')$ does not necessarily mean that these are congruent: one tetrahedron could be the mirror image of the other. This test eliminates 2 of the remaining 4 solutions. Finally, testing for “non-rigidity” shows that the remaining 2 solutions are indeed rigid. When the parts are assembled to form one packing, holding spheres 1,2,3 in place, these two solutions are mirror images through the plane of spheres 1,2,3. However, as can be seen in Figure 3, the packing is achiral (mirror self-symmetric), so these two solutions are both instances of the same arrangement. In summary, this particular set of contact pairs for $N = 12$ has exactly 1 realizable and rigid sphere packing, which is presented in Figure 3. For better visibility, the same packing is also shown using smaller spheres with the addition of rods that indicate the pairwise contacts. Spheres 5, 6, 7, and 8 have been darkened in the figure to show how the full assembly was cut into two pieces for solving incrementally.

We now return to computing the set C_8 , with the following using only serial computations. Rather than computing it from scratch all in one blow as above, we could have computed this by adding new vertices one at a time. For example, suppose we start with the curve C_4 of degree 2 corresponding to four vertices v_1, \dots, v_4 with v_1, \dots, v_3 fixed as above and contact pairs $\{1-2, 1-3, 1-4, 2-3, 2-4\}$.

Building on C_4 , we add vertex v_5 with new contact pairs $\{1-5, 3-5, 4-5\}$ yielding

$$(C_4 \times \mathbb{C}^3) \cap \mathcal{V}(D_{1,5}, D_{3,5} - D_{1,5}, D_{4,5}).$$

In under a second, Algorithm 1 computes the irreducible decomposition as follows:

- one-dimensional component of degree 2 with $v_2 = v_5$;
- one-dimensional component of degree 2 with $v_3 = v_4$;
- one-dimensional component, call it C_5 , of degree 4.

Next, we add vertex v_6 with new contact pairs $\{2-6, 3-6, 4-6\}$ by computing

$$(C_5 \times \mathbb{C}^3) \cap \mathcal{V}(D_{2,6}, D_{3,6} - D_{2,6}, D_{4,6}).$$

In 2 seconds, Algorithm 1 computes the irreducible decomposition as follows:

- one-dimensional component of degree 4 with $v_6 = v_1$;
- one-dimensional component, call it C_6 , of degree 4.

Adding vertex v_7 with new contact pairs $\{3-7, 5-7, 6-7\}$ is described via

$$(C_6 \times \mathbb{C}^3) \cap \mathcal{V}(D_{3,7}, D_{5,7}, D_{6,7}).$$

In 5 seconds, Algorithm 1 computes the irreducible decomposition as follows:

- two one-dimensional components of degree 2 with $v_5 = v_6$;
- one-dimensional component, call it C_7 , of degree 12.

Finally, we add vertex v_8 with new contact pairs $\{4-8, 5-8, 6-8\}$ by computing

$$(C_7 \times \mathbb{C}^3) \cap \mathcal{V}(D_{4,8}, D_{5,8}, D_{6,8}).$$

In 18 seconds, Algorithm 1 computes the irreducible components are A_1 and A_2 which are described above. Thus, this sequence of solves computes $C_8 = A_1 \cup A_2$ in approximately 25 seconds using only serial processing.

Not only is the total computation time to find and decompose C_8 in a sequence of intersections less than that used by the regenerative cascade applied to the whole system for C_8 , the sequence of intersections has also generated witness sets for C_4 , C_5 , C_6 , and C_7 along the way. These can be re-used in other intersections if one seeks to more thoroughly explore the full set of sphere packings.




Figure 3: An $N = 12$ sphere packing showing (a) spheres in contact, and (b) smaller spheres with rods indicating contact pairs. Dark spheres indicate the loop 5-8-6-7-5.

9 Conclusions

Even though Lemma 4.1, Theorem 5.1, and Theorem 6.2 are proven using modest modifications of existing results in diagonal intersection and isosingular theory, together they bring about a substantial new flexibility in how one may numerically solve polynomial systems and numerically explore questions in algebraic geometry. Instead of just aiming at the solution set of a final system, one is free to design a sequence of intersection operations that arrive at the final result. In our final example, a few inexpensive operations to find the numerical irreducible decomposition of an 8-sphere arrangement allows the final intersections to solve a related 12-sphere arrangement to be done by tracking just $(12^2 + 12 \cdot 14 + 14^2) \cdot 2 \cdot 2 = 2032$ paths instead of the 16,777,216 paths that a total-degree or polyhedral homotopy would require. Moreover, along the way, one obtains solutions for several smaller arrangements of spheres that could be combined in alternative ways to arrive at other large packings. Similarly, the four-bar path-synthesis example generated witness sets for all the two- to eight-point problems on the way to the final nine-point problem. As the reported run times attest, the new approach is faster than existing methods.

After one finds a witness set for an algebraic set of interest, the new method allows one to impose new conditions that pick out subsets for further study. These new conditions may involve new variables, such as when one formulates singularity conditions in terms of tangent vectors. This *extension* capability provides a kind of flexibility that is compatible with the way mathematicians, engineers, and scientists tend to approach problem areas.

Finally, the diagonal intersection technique of [42] and the homotopy membership test of [47] are special cases of the new method. This unification is useful for designing streamlined software.

10 Acknowledgment

We would like to thank Michael Brenner of the Harvard SEAS Department for discussions on sphere packings and for suggesting the subdivide-and-merge approach to solving them. This helped motivate the formulation of our new approach. We also would like to thank the referees for their many helpful comments.

References

- [1] H. Alt, Über die erzeugung gegebener ebener kurven mit hilfe des gelenkvierecks, *ZAMM*, 3(1), 13–19, 1923.
- [2] N. Arkus, V.N. Manoharan, and M.P. Brenner, Deriving finite sphere packings. *SIAM J. Discrete Math* 25(4), 1860–1901, 2011.
- [3] V.I. Arnol’d. Singularities of smooth mappings. *Russian Math. Surveys* 23, 1–43, 1968.
- [4] P. Aubry, F. Rouillier, and M. Safey El Din. Real solving for positive dimensional systems. *J. Symbolic Comput.*, 34(6), 543–560, 2002.
- [5] L. Blum, F. Cucker, M. Shub, and S. Smale. *Complexity and real computation*, Springer-Verlag, New York, 1998.

- [6] D.J. Bates, J.D. Hauenstein, C. Peterson, and A.J. Sommese, Numerical decomposition of the rank-deficiency set of a matrix of multivariate polynomials. In *Approximate Commutative Algebra*, vol. 14 of *Texts and Monographs in Symbolic Computation*, Springer, 2010, pp. 55–77.
- [7] D.J. Bates, J.D. Hauenstein, C. Peterson, and A.J. Sommese, A numerical local dimension test for points on the solution set of a system of polynomial equations. *SIAM J. Num. Anal.* 47(5), 3608–3623, 2009.
- [8] D.J. Bates, J.D. Hauenstein, A.J. Sommese, and C.W. Wampler, Bertini: Software for numerical algebraic geometry. Available at bertini.nd.edu.
- [9] D.J. Bates, J.D. Hauenstein, A.J. Sommese, and C.W. Wampler, *Numerically Solving Polynomial Systems with Bertini*. SIAM, 2013.
- [10] D.J. Bates, C. Peterson, A.J. Sommese, and C.W. Wampler, Numerical computation of the genus of an irreducible curve within an algebraic set, *J. Pure Appl. Algebra* 215, 1844–1851, 2011.
- [11] G.M. Besana, S. Di Rocco, J.D. Hauenstein, A.J. Sommese, and C.W. Wampler, Cell decomposition of almost smooth real algebraic surfaces. *Numer. Algorithms*, 63(4), 645–678, 2013.
- [12] J. Boardman, Singularities of differentiable maps. *Inst. Hautes Études Sci. Publ. Math.* 33, 21–57, 1967.
- [13] B. Dayton and Z. Zeng. Computing the multiplicity structure in solving polynomial systems. In *ISSAC'05*, ACM, New York, 2005, pp. 116–123.
- [14] J. Draisma, E. Horobet, G. Ottaviani, B. Sturmfels, and R.R. Thomas. The Euclidean distance degree of an algebraic variety. In *SNC '14*, ACM, New York, 2014, pp. 9–16.
- [15] J. Harris, *Algebraic geometry: a first course*, Springer-Verlag, New York, 1992.
- [16] B. Hassett. *Introduction to Algebraic Geometry*, Cambridge University Press, Cambridge, 2007.
- [17] J.D. Hauenstein, Numerically computing real points on algebraic sets. *Acta. Appl. Math.* 125(1), 105–119, 2013.
- [18] J.D. Hauenstein and A.J. Sommese, Membership tests for images of algebraic sets by linear projections. *Appl. Math. Comput.* 219(12), 6809–6818, 2013.
- [19] J.D. Hauenstein and A.J. Sommese, Witness sets of projections. *Appl. Math. Comput.* 217(7), 3349–3354, 2010.
- [20] J.D. Hauenstein, A.J. Sommese, and C.W. Wampler, Regeneration homotopies for solving systems of polynomials. *Math. Comp.* 80, 345–377, 2011.
- [21] J.D. Hauenstein, A.J. Sommese, and C.W. Wampler, Regenerative cascade homotopies for solving polynomial systems. *Appl. Math. Comput.* 218(4), 1240–1246, 2011.
- [22] J.D. Hauenstein and C.W. Wampler, Isosingular sets and deflation. *Found. Comput. Math.* 13(3), 371–403, 2013.

- [23] J.D. Hauenstein, and C.W. Wampler, Numerically intersecting algebraic varieties via witness sets. *Appl. Math. Comput.* 219(10), 5730–5742, 2013.
- [24] S. Hormoz and M.P. Brenner, Design principles for self-assembly with short-range interactions. *Proc. National Academy of Sciences* 108(13), 5193–5198, 2011.
- [25] B. Huber and B. Sturmfels, A polyhedral method for solving sparse polynomial systems. *Math. Comp.*, 64(212), 1541–1555, 1995.
- [26] B. Huber and J. Verschelde. Polyhedral end games for polynomial continuation. *Numer. Algorithms*, 18(1), 91–108, 1998.
- [27] G. Kemper. *A Course in Commutative Algebra*, Springer, Heidelberg, 2011.
- [28] T.-L. Lee, T.Y. Li, and C.H. Tsai. HOM4PS-2.0: a software package for solving polynomial systems by the polyhedral homotopy continuation method. *Computing*, 83, 109–133, 2008.
- [29] A. Leykin. Numerical primary decomposition. In *ISSAC 2008*, ACM, New York, 2008, pp. 165–172.
- [30] A. Leykin, J. Verschelde, and A. Zhao. Newton’s method with deflation for isolated singularities of polynomial systems. *Theoret. Comput. Sci.*, 359(1-3), 111–122, 2006.
- [31] T.Y. Li, Numerical solution of polynomial systems by homotopy continuation methods, in *Handbook of numerical analysis, Vol. XI*, North-Holland, 209–304, 2003.
- [32] Y. Lu, D.J. Bates, A.J. Sommese, and C.W. Wampler, Finding all real points of a complex curve. *Contemp. Math.* 448, 183–205, 2007.
- [33] A. P. Morgan, *Solving Polynomial Systems Using Continuation for Engineering and Scientific Problems*. Prentice Hall, 1987.
- [34] A.P. Morgan, and A.J. Sommese, A homotopy for solving general polynomial systems that respects m -homogeneous structures. *Appl. Math. Comput.* 24(2), 101–113, 1987.
- [35] A.P. Morgan and A.J. Sommese, Coefficient-parameter polynomial continuation. *Appl. Math. Comput.* 29(2), 123–160, 1989. Errata: *Appl. Math. Comput.*, 51(2-3), 207, 1992.
- [36] A.P. Morgan, A.J. Sommese, L.T. Watson. Finding all isolated solutions to polynomial systems using HOMPACT. *ACM ToMS*, 15(2), 93–122, 1989.
- [37] A.P. Morgan, A.J. Sommese, and C.W. Wampler. A power series method for computing singular solutions to nonlinear analytic systems. *Numerische Math.*, 63, 391–409, 1992.
- [38] A.P. Morgan, A.J. Sommese, and C.W. Wampler. Computing singular solutions to polynomial systems. *Adv. in Appl. Math.*, 13, 305–327, 1992.
- [39] K. Ranestad and B. Sturmfels, The convex hull of a variety. In *Notions of Positivity and the Geometry of Polynomials*, a volume of *Trends in Mathematics*, Springer Basel, 2011, pp. 331–344.

- [40] F. Rouillier, M.-F. Roy, and M. Safey El Din. Finding at least one point in each connected component of a real algebraic set defined by a single equation. *J. Complexity*, 16(4), 716–750, 2000.
- [41] S. M. Rump, Verification methods: Rigorous results using floating point arithmetic. *Acta Numerica*, 19, 287–449, 2010.
- [42] A.J. Sommese, J. Verschelde, and C.W. Wampler, Homotopies for intersecting solution components of polynomial systems. *SIAM J. Numer. Anal.* 42(4), 1552–1571, 2004.
- [43] A.J. Sommese, J. Verschelde, and C.W. Wampler, An intrinsic homotopy for intersecting algebraic varieties. *J. Complexity* 21(4), 593–608, 2005.
- [44] A.J. Sommese, J. Verschelde, and C.W. Wampler, Numerical decomposition of the solution sets of polynomial systems into irreducible components. *SIAM J. Numer. Anal.* 38(6), 2022–2046, 2001.
- [45] A.J. Sommese, J. Verschelde, and C.W. Wampler, Numerical irreducible decomposition using projections from points on the components. *Contemp. Math.* 286, 37–51, 2001.
- [46] A.J. Sommese, J. Verschelde, and C.W. Wampler, Symmetric functions applied to decomposing solution sets of polynomial systems. *SIAM J. Numer. Anal.* 40(6), 2026–2046, 2002.
- [47] A.J. Sommese, J. Verschelde, and C.W. Wampler, Using monodromy to decompose solution sets of polynomial systems into irreducible components. In *Applications of algebraic geometry to coding theory, physics and computation (Eilat, 2001)*, vol. 36 of *NATO Sci. Ser. II Math. Phys. Chem.*, Kluwer Acad. Publ., 2001, pp. 297–315.
- [48] A.J. Sommese, J. Verschelde, and C.W. Wampler, Solving polynomial systems equation by equation, in *Algorithms in algebraic geometry*, IMA Vol. Math. Appl., 146, 133–152, Springer, 2008.
- [49] A.J. Sommese and J. Verschelde, Numerical homotopies to compute generic points on positive dimensional algebraic sets. *J. Complexity* 16(3), 572–602, 2000.
- [50] A.J. Sommese and C.W. Wampler, Numerical algebraic geometry, in *The mathematics of numerical analysis (Park City, UT, 1995)*. *Lectures Appl. Math.*, 32, 749–763, AMS, 1996.
- [51] A.J. Sommese and C.W. Wampler, *The Numerical Solution of Systems of Polynomials Arising in Engineering and Science*, World Scientific, Singapore, 2005.
- [52] A.J. Sommese and C.W. Wampler, Exceptional sets and fiber products. *Found. Comp. Math.*, 28, 171–196, 2008.
- [53] H.-J. Su, J.M. McCarthy, M. Sosonkina and L.T. Watson. Algorithm 857: POLSYS_GLP—a parallel general linear product homotopy code for solving polynomial systems of equations. *ACM ToMS*, 32, 561–579, 2006.
- [54] L.N. Trefethen and J.A.C. Weideman, The exponentially convergent trapezoidal rule. *SIAM Review*, 56(3), 385–458, 2014.

- [55] J. Verschelde, Algorithm 795: PHCpack: A general-purpose solver for polynomial systems by homotopy continuation. *ACM ToMS*, 25(2), 251–276, 1999.
- [56] J. Verschelde and R. Cools, Symbolic homotopy construction. *Appl. Alg. Engrg. Comm. Comput.*, 4(3), 169–183, 1993.
- [57] J. Verschelde, P. Verlinden, and R. Cools, Homotopies exploiting newton polytopes for solving sparse polynomial systems. *SIAM J. Numer. Anal.*, 31(3), 915–930, 1994.
- [58] C.W. Wampler, A.P. Morgan, and A.J. Sommese, Complete solution of the nine-point path synthesis problem for four-bar linkages. *ASME J. Mech. Des.*, 114, 153–159, 1992.