



HAL
open science

A General Solver Based on Sparse Resultants: Numerical Issues and Kinematic Applications

Ioannis Z. Emiris

► **To cite this version:**

Ioannis Z. Emiris. A General Solver Based on Sparse Resultants: Numerical Issues and Kinematic Applications. RR-3110, INRIA. 1997. inria-00073580

HAL Id: inria-00073580

<https://inria.hal.science/inria-00073580>

Submitted on 24 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

***A General Solver Based on Sparse Resultants:
Numerical Issues and Kinematic Applications***

Ioannis Z. Emiris

N° 3110

January 1997

————— THÈME 2 —————



*rapport
de recherche*



A General Solver Based on Sparse Resultants: Numerical Issues and Kinematic Applications

Ioannis Z. Emiris *

Thème 2 — Génie logiciel
et calcul symbolique
Projet SAFIR

Rapport de recherche n° 3110 — January 1997 — 29 pages

Abstract: Sparse elimination and the sparse resultant exploit the structure of polynomials by measuring their complexity in terms of Newton polytopes instead of total degree. We sketch the sparse resultant constructions of Canny and Emiris and show how they reduce the problem of root-finding to an eigenproblem. A little known method for achieving this reduction is presented which does not increase the dimension of the problem. Together with an implementation of the sparse resultant construction, this provides a general solver for polynomial systems. We discuss the overall implementation and emphasize the numerical issues encountered in the matrix computations, based on the condition number of a matrix. We illustrate the program's capabilities by applying it to concrete problems from vision, robotics and computational biology. The high efficiency and accuracy of the solutions suggest that sparse elimination may be the method of choice for systems of moderate size.

Key-words: polynomial equation, system solving, sparse resultant, eigenvalue and eigenvector, matrix condition number, vision, robotics, computational biology

(Résumé : tsvp)

* INRIA, B.P. 93, Sophia-Antipolis 06902, France emiris@sophia.inria.fr.

Un Logiciel Général pour la Résolution des Systèmes Algébriques Basé sur le Résultant Creux : Aspects Numériques et Applications en Cinématique

Résumé : La théorie d'élimination creuse et le résultant creux exploitent la structure des polynômes en regardant leur complexité en termes des polytopes de Newton au lieu du degré total. Nous décrivons en gros les constructions du résultant creux proposées par Canny et Emiris, et nous montrons comment elles réduisent la résolution d'un système au calcul de valeurs et vecteurs propres. Une méthode peu connue pour cette réduction est présentée, et qui n'augmente pas la dimension du problème. Nous discutons l'implémentation générale en mettant l'accent sur les questions numériques rencontrées en le calcul matriciel à partir du conditionnement d'une matrice. Les fonctionnalités du logiciel sont démontrées par des applications en vision, en robotique et en biologie algorithmique. L'efficacité du programme et la précision des solutions indiquent que l'élimination creuse serait la méthode préférée pour des systèmes de taille moyenne.

Mots-clé : équation polynomiale, résolution de systèmes, résultant creux, valeurs et vecteurs propres, conditionnement de matrices, vision, robotique, biologie algorithmique

1 Introduction

The problem of computing all common zeros of a system of polynomials is of fundamental importance in a wide variety of scientific and engineering applications. This article surveys an efficient method based on the sparse resultant for computing all *isolated* solutions of an arbitrary system of n polynomials in n unknowns. In particular, we exploit the algorithms of Canny and Emiris [CE93, EC95] for constructing sparse resultant formulae which yield nontrivial multiples of the resultant. We show that the matrices obtained allow the reduction of the root-finding problem to the eigendecomposition of a square matrix. The emphasis here is placed on numerical issues and the application of our implementation to concrete problems.

We describe very briefly the main steps in sparse elimination and the construction of sparse resultant matrices. Most proofs are omitted but can be found in [CE93, Emi94, EC95, Emi96]. The study of coordinate rings of varieties in K^n , where K is a field, has been shown to be particularly useful in studying systems of polynomial equations. We concentrate on zero-dimensional varieties for which it is known that the coordinate ring forms a finite-dimensional vector space and, actually, an algebra over K . An important algorithmic question is the construction of an explicit monomial K -basis for such a space. Based on monomial bases, we may generate generic endomorphisms or multiplication maps for any given polynomial, as outlined in section 2.

Root finding is reduced to an eigenproblem and then existing techniques are employed from numerical linear algebra. An important feature of our method is precisely that it reduces to matrix operations for which powerful and accurate implementations already exist. Section 3 discusses this method for the case of adding an extra u -polynomial to obtain an overconstrained system. This is the classical method, used in defining the u -resultant; it possesses the advantage that the matrices have a lot of known structure.

A relatively novel approach that keeps the number of polynomials fixed is proposed in section 4 where one of the variables is *hidden* in the coefficient field, thus producing an overconstrained system. We show how the calculation of all isolated roots again reduces to an eigenproblem. Keeping the number of polynomials fixed has been observed to be an important advantage in practice. On the other hand, it leads to arbitrary matrix polynomials for which we have to calculate all eigenvalues and eigenvectors.

This approach has been implemented in C and provides, together with an eigenvalue solver, a self-contained and fast polynomial solver. Section 5 describes our

implementation and discusses the practical issues that arise thereof. In particular, we consider issues of numerical stability and conditioning of the matrices.

Our techniques find their natural application in problems arising in a variety of fields, including problems expressed in terms of geometric and kinematic constraints. As an empirical observation, polynomial systems encountered in robot and molecular kinematics, motion and aspect ratio calculation in vision and geometric modeling are characterized by small mixed volume compared to their Bézout bound. The complexity of our methods depends directly on this sparse structure, in contrast to Gröbner bases. Sparse homotopies were proposed in order to exploit the same structure [HS95, VVC94], yet they may suffer from the usual accuracy and completeness problems of numerical techniques. Lastly, resultant-based methods include a large fraction of offline computation: the first phase, including the construction of the matrix, has to be executed only once for a given set of supports. For every specific instance, the coefficients are specialized and then the online phase takes place.

We describe in detail two problems from vision, robot kinematics and structural biology. The first problem, analyzed and solved in section 6, is a standard problem from photogrammetry. Given information about a static scene seen by two positions of a camera, the camera motion must be computed. When the minimum amount of information is available so that the problem is solvable, an optimal sparse resultant matrix is constructed and the numerical answers are computed efficiently and accurately by our implementation. Our method exhibits competitive speed, as compared to previous approaches, and better accuracy.

The second application, in section 7, comes from computational biology and reduces to an inverse kinematics problem. The symmetry of the molecule at hand explains the high multiplicity of the common zeros, which leads us to compare the two approaches of defining an overconstrained system, either by adding a u -polynomial or by hiding one of the input variables. Both methods are used for different instances in order for all roots to be calculated accurately.

We conclude with some open questions in section 8. Related work is mentioned in each corresponding section.

2 Sparse Elimination

Sparse elimination generalizes several results of classical elimination theory on multivariate polynomial systems of arbitrary degree by considering the structure of the given polynomials, namely their Newton polytopes. This leads to stronger algebraic and combinatorial results in general. Assume that the number of variables is n ; roots

in $(\mathbb{C}^*)^n$ are called *toric*. By concentrating on \mathbb{C}^* we may, consequently, extend our scope to *Laurent polynomials*. We use x^e to denote the monomial $x_1^{e_1} \cdots x_n^{e_n}$, where $e = (e_1, \dots, e_n) \in \mathbb{Z}^n$ is an exponent vector or, equivalently, an integer lattice point, and $n \in \mathbb{Z}_{\geq 1}$. Let the input Laurent polynomials be

$$f_1, \dots, f_n \in K[x_1, x_1^{-1}, \dots, x_n, x_n^{-1}] = K[x, x^{-1}] \quad (1)$$

where K is a field.

Let $\mathcal{A}_i = \text{supp}(f_i) = \{a_{i1}, \dots, a_{i\mu_i}\} \subset \mathbb{Z}^n$ denote the set, with cardinality μ_i , of exponent vectors corresponding to monomials in f_i with nonzero coefficients. This set is the *support* of f_i :

$$f_i = \sum_{a_{ij} \in \mathcal{A}_i} c_{ij} x^{a_{ij}}, \quad c_{ij} \neq 0.$$

Definition 2.1 *The Newton polytope of f_i is the convex hull of support \mathcal{A}_i , denoted $Q_i = \text{Conv}(\mathcal{A}_i) \subset \mathbb{R}^n$.*

For arbitrary sets in \mathbb{R}^n there is a natural associative and commutative addition operation called Minkowski addition.

Definition 2.2 *The Minkowski sum $A + B$ of sets A and B in \mathbb{R}^n is*

$$A + B = \{a + b \mid a \in A, b \in B\} \subset \mathbb{R}^n.$$

If A and B are convex polytopes then $A + B$ is a convex polytope.

Let $\text{Vol}(A)$ denote the standard (Lebesgue) volume of polytope A in n -dimensional euclidean space.

Definition 2.3 *Given convex polytopes $A_1, \dots, A_n \subset \mathbb{R}^n$, there is a unique real-valued function $MV(A_1, \dots, A_n)$, called the mixed volume of A_1, \dots, A_n which has the following two properties. First, it is multilinear with respect to Minkowski addition and scalar multiplication i.e. for $\mu, \rho \in \mathbb{R}_{\geq 0}$ and convex polytope $A'_k \subset \mathbb{R}^n$*

$$\begin{aligned} & MV(A_1, \dots, \mu A_k + \rho A'_k, \dots, A_n) = \\ & = \mu MV(A_1, \dots, A_k, \dots, A_n) + \rho MV(A_1, \dots, A'_k, \dots, A_n). \end{aligned}$$

Second,

$$MV(A_1, \dots, A_n) = n! \text{Vol}(A_1), \quad \text{when } A_1 = \dots = A_n.$$

Notationally, we use

$$MV(Q_1, \dots, Q_n) = MV(\mathcal{A}_1, \dots, \mathcal{A}_n) = MV(f_1, \dots, f_n).$$

We are now ready to state Bernstein's theorem [Ber75], the cornerstone of sparse elimination, generalized to arbitrary varieties.

Theorem 2.4 [Ful93, sect. 5.5] *Given are polynomials $f_1, \dots, f_n \in K[x, x^{-1}]$ with Newton polytopes Q_1, \dots, Q_n . For any isolated common zero $\alpha \in (\mathbb{C}^*)^n$, let $i(\alpha)$ denote the intersection multiplicity at this point. Then $\sum_{\alpha} i(\alpha) \leq MV(Q_1, \dots, Q_n)$, where the sum ranges over all isolated roots. Equality holds when all coefficients are generic.*

Canny and Rojas have substantially weakened the requirements for equality [CR91]. A recent result generalizes the notion of mixed volume to that of *stable volume*, thus extending the bound to non-toric roots.

Theorem 2.5 [HS96] *For polynomials $f_1, \dots, f_n \in \mathbb{C}[x, x^{-1}]$ with supports $\mathcal{A}_1, \dots, \mathcal{A}_n$ the number of common isolated zeros in \mathbb{C}^n , counting multiplicities, is upwards bounded by the stable volume of $\mathcal{A}_1, \dots, \mathcal{A}_n$.*

Bernstein's bound is at most as high as Bézout's bound, which is simply the product of the total degrees, and is usually significantly smaller for systems encountered in engineering applications.

The *sparse* or *Newton resultant* provides a necessary and generically sufficient condition for the existence of toric roots for a system of $n + 1$ polynomials in n variables:

$$f_1, \dots, f_{n+1} \in K[x, x^{-1}]. \quad (2)$$

To define the sparse resultant we regard a polynomial f_i as a generic point $c_i = (c_{i1}, \dots, c_{im_i})$ in the space of all possible polynomials with the given support $\mathcal{A}_i = \text{supp}(f_i)$, where m_i is the number of nonzero terms. It is natural to identify scalar multiples, so the space of all such polynomials contracts to the projective space $\mathbb{P}_K^{m_i-1}$ or, simply, \mathbb{P}^{m_i-1} . Then the input system (2) can be thought of as a point

$$c = (c_1, \dots, c_{n+1}) \in \mathbb{P}^{m_1-1} \times \dots \times \mathbb{P}^{m_{n+1}-1}.$$

Let $Z_0 = Z_0(\mathcal{A}_1, \dots, \mathcal{A}_{n+1})$ be the set of all points c such that the system has a solution in $(\mathbb{C}^*)^n$ and let $Z = Z(\mathcal{A}_1, \dots, \mathcal{A}_{n+1})$ denote the Zariski closure of Z_0 in the product of projective spaces. It is proven in [PS93] that Z is an irreducible variety.

Definition 2.6 *The sparse resultant $R = R(\mathcal{A}_1, \dots, \mathcal{A}_{n+1})$ of system (2) is a polynomial in $\mathbb{Z}[c]$. If $\text{codim}(Z) = 1$ then $R(\mathcal{A}_1, \dots, \mathcal{A}_{n+1})$ is the defining irreducible polynomial of the hypersurface Z . If $\text{codim}(Z) > 1$ then $R(\mathcal{A}_1, \dots, \mathcal{A}_{n+1}) = 1$.*

Let $\deg_{f_i} R$ denote the degree of the resultant R in the coefficients of polynomial f_i and let

$$MV_{-i} = MV(Q_1, \dots, Q_{i-1}, Q_{i+1}, \dots, Q_{n+1}) \quad \text{for } i = 1, \dots, n+1.$$

A consequence of Bernstein's theorem is

Theorem 2.7 [PS93] *The sparse resultant is separately homogeneous in the coefficients c_i of each f_i and its degree in these coefficients equals the mixed volume of the other n Newton polytopes i.e. $\deg_{f_i} R = MV_{-i}$.*

Canny and Emiris [CE93, EC95] have proposed the first two efficient algorithms for constructing *sparse resultant*, or *Newton matrices*, i.e., matrices in the coefficients whose determinant is a nontrivial multiple of the sparse resultant. The first algorithm relies on a mixed subdivision of the Minkowski Sum, while the second constructs the matrix in an incremental fashion. For those cases where it is provably possible, the incremental algorithm yields optimal matrices, so that the determinant equals the resultant. For general systems, it typically produces matrices that are at most 3 times larger than optimal. Let

$$Q = Q_1 + \dots + Q_{n+1} \subset \mathbb{R}^n$$

be the Minkowski sum of all Newton polytopes. Let the resultant matrix be M . The subdivision-based algorithm lets

$$\mathcal{E} = (Q + \delta) \cap \mathbb{Z}^n$$

be the set that indexes the rows and columns of M in a bijective way, where $\delta \in \mathbb{Q}^n$ is an arbitrarily small and *sufficiently generic* vector. Clearly, M has dimension $|\mathcal{E}|$. The incremental algorithm also indexes the rows and columns with monomials, or equivalently, lattice points which lie in \mathcal{E} . This algorithm, though, is different and produces significantly smaller matrices but we have no formal result on their dimension besides the fact that the matrix dimension is bounded by that of the subdivision-based construction. Irrespective of the algorithm applied, the resultant matrix M has the following properties.

Theorem 2.8 [Emi94] *Matrix M obtained by either algorithm is well-defined, square, generically nonsingular and its determinant is divisible by the sparse resultant R .*

To be more precise, the rows of M are indexed by a pair composed of a monomial and an input polynomial. The entries of the respective row are coefficients of this polynomial. The degree of $\det M$ in the coefficients of f_i for $i = 1, \dots, n+1$ is greater or equal to MV_{-i} .

To solve system (1) we define an overconstrained system by one of the two ways below. We apply the resultant matrix construction on the new system and use the following properties. Let $I = I(f_1, \dots, f_n)$ be the ideal generated by polynomials (1) and $V = V(f_1, \dots, f_n) \in (\overline{K}^*)^n$ their variety, where \overline{K} is the algebraic closure of field K . We are interested in the case that V has *dimension zero*. Then, its coordinate ring $K[x, x^{-1}]/I$ is an m -dimensional vector space over K by theorem 2.4, where

$$m = MV(f_1, \dots, f_n) = MV(Q_1, \dots, Q_n).$$

Using the subdivision-based construction it is easy to show [ER94, Emi96] that generically a monomial basis of $K[x, x^{-1}]/I$ can be found among the monomials of

$$Q_1 + \dots + Q_n \subset \mathbb{R}^n.$$

Moreover, resultant matrix M produces the *multiplication map* for any given f_0 . This is a matrix, or an endomorphism, that serves in computing in the coordinate ring and essentially allows computation of the common roots of $f_1 = \dots = f_n = 0$.

3 Adding a Polynomial

The problem addressed here is to find all isolated roots $\alpha \in V$ where $V \subset (\overline{K}^*)^n$ is the zero-dimensional variety of (1), with cardinality bounded by $m = MV(Q_1, \dots, Q_n)$. In addition to zero-dimensional, the ideal $I = I(f_1, \dots, f_n)$ is assumed to be *radical*, or self-radical i.e., $I = \sqrt{I}$, which is equivalent to saying that all roots in V are distinct. This requirement is weakened later.

An overconstrained system is obtained by adding extra polynomial f_0 to the given system. We choose f_0 to be linear with coefficients c_{0j} and constant term equal to indeterminate u .

$$f_0 = u + c_{01}x_1 + \dots + c_{0n}x_n \in (K[u])[x, x^{-1}].$$

Coefficients c_{0j} , $j = 1, \dots, n$, should define an *injective* function

$$f_0 : V \rightarrow \overline{K} : \alpha \mapsto f_0(\alpha).$$

There are standard deterministic strategies for selecting c_{0j} so that they ensure the injective property. In practice, coefficients c_{0j} may be randomly distributed in some range of integer values of size $S > 1$, and a bad choice for c_{01}, \dots, c_{0n} is one that will result in the same value of $f_0 - u$ at two distinct roots α and α' . Assume that α and α' differ in their i -th coordinate for some $i > 0$, then fix all choices of c_{0j} for $j \neq i$; the probability of a bad choice for c_{0i} is $1/S$, and since there are $\binom{m}{2}$ pairs of roots, the total probability of failure for this scheme is

$$\text{Prob}[\text{failure}] \leq \binom{m}{2}/S : \quad c_{0j} \in \{1, \dots, S\}, j = 1, \dots, n.$$

It suffices, therefore, to pick c_{0j} from a sufficiently large range in order to make the probability of success arbitrarily high. Moreover, it is clear that any choice of coefficients can be tested deterministically at the end of the algorithm.

Either algorithm for the resultant matrix may be used to build matrix M . As before, the vanishing of $\det M$ is a necessary condition for the overconstrained system to have common roots. For $\alpha \in V$, u is constrained to a specific value determined by the c_{0j} coefficients. The construction of M is not affected by this definition of f_0 . Let monomial set \mathcal{E} index the columns of M and partition M so that the lower right square submatrix M_{22} depends on u and has size r . Suppose for now that the upper left square submatrix M_{11} is *nonsingular*. Clearly $r \geq m$ and equality holds if M is obtained by the subdivision resultant algorithm. By the construction of M and for an arbitrary $\alpha \in (\overline{K}^*)^n$

$$\begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22}(u) \end{bmatrix} \begin{bmatrix} \vdots \\ \alpha^q \\ \vdots \end{bmatrix} = \begin{bmatrix} \vdots \\ \alpha^p f_{i_p}(\alpha) \\ \vdots \end{bmatrix} : \quad q, p \in \mathcal{E}, i_p \in \{0, 1, \dots, n\}, \quad (3)$$

Now $M'(u) = M_{22}(u) - M_{21}M_{11}^{-1}M_{12}$ where its diagonal entries are linear polynomials in u and

$$M'(u) \begin{bmatrix} \alpha^{b_1} \\ \vdots \\ \alpha^{b_r} \end{bmatrix} = \begin{bmatrix} \alpha^{p_1} f_0(\alpha, u) \\ \vdots \\ \alpha^{p_r} f_0(\alpha, u) \end{bmatrix}, \quad (4)$$

where $\mathcal{B} = \{b_1, \dots, b_r\}$ and $\{p_1, \dots, p_r\}$ index the columns and rows, respectively, of M_{22} and thus M' . For a root $\alpha \in V$ and for

$$u = - \sum_{j=1}^n c_{0j} \alpha_j$$

the right hand side vector in (4) is null. Let $v'_\alpha = [\alpha^{b_1}, \dots, \alpha^{b_r}]$ and write $M'(u) = M' + uI$, where M' now is numeric and I is the $r \times r$ identity matrix. Then

$$(M' + uI)v'_\alpha = 0 \Rightarrow \left[M' - \left(\sum_j c_{0j} \alpha_{ij} \right) I \right] v'_\alpha = 0.$$

This essentially reduces root-finding to an eigenproblem since, for every solution of the original system, there is an eigenvalue and eigenvector of M and hence of M' . Below we study how to compute a candidate solution from every eigenvalue-eigenvector pair.

If the generated ideal I is radical then every eigenvalue has *algebraic multiplicity* one with probability greater or equal to $1 - \binom{m}{2}/S$. We can weaken the condition that I be radical by requiring only that each eigenvalue has *geometric multiplicity* one. This equals the dimension of the eigenspace associated with an eigenvalue. If there exist eigenvalues of higher geometric multiplicity this technique fails. Several remedies exist, including the definition of an overconstrained system by hiding a variable, as in the next section, and then deriving the root coordinates one by one.

In what follows we assume that all eigenvalues have unit geometric multiplicity. Hence it is guaranteed that among the eigenvectors of M' we shall find the vectors v'_α for $\alpha \in V$. By construction of M [Emi96] each eigenvector v'_α of M' contains the values of monomials \mathcal{B} at some common root $\alpha \in (\overline{K}^*)^n$. By (3) we can define vector v_α as follows:

$$M_{11}v_\alpha + M_{12}v'_\alpha = 0 \Rightarrow v_\alpha = -M_{11}^{-1}M_{12}v'_\alpha. \quad (5)$$

The size of v_α is $|\mathcal{E}| - r$, indexed by $\mathcal{E} \setminus \mathcal{B}$. It follows that vectors v_α and v'_α together contain the values of every monomial in \mathcal{E} at some root α .

Theorem 3.1 [Emi96] *Assume \mathcal{E} spans \mathbb{Z}^n . Then there exists a polynomial-time algorithm that finds a subset of $n + 1$ affinely independent points in \mathcal{E} . Given v_α , v'_α and these points, we can compute the coordinates of root $\alpha \in V(I)$. If all $n + 1$ independent points are in \mathcal{B} then v'_α suffices.*

In practice, most of the operations described here are not implemented with (exact) rational arithmetic but are instead carried out over floating point numbers of fixed size. An important aspect of this computation is numerical error, which we discuss below in a separate section and in the particular context of specific applications later.

Let $MM(k)$ express the time complexity for multiplying two $k \times k$ matrices, currently at $\mathcal{O}(n^{2.376})$ [CW90].

Theorem 3.2 [Emi96] *Suppose that system (1) generates a zero-dimensional radical ideal, matrix M has been computed such that M_{11} is nonsingular and $n+1$ affinely independent points in \mathcal{E} have been computed. Then all common zeros of the polynomial system are approximated in time*

$$\mathcal{O}(MM(|\mathcal{E}|) + |\mathcal{E}|n^2 + MM(n)m).$$

It is clear that for most systems the arithmetic complexity is dominated by the first term, namely the complexity of matrix operations and in particular the eigen-decomposition. Under reasonable assumptions the complexity becomes [Emi96]

$$e^{\mathcal{O}(n)} m^3.$$

As expected, the complexity is single exponential in n and polynomial in the number of roots.

One hypothesis concerned the nonsingularity of M_{11} . When it is singular the resultant matrix is regarded as a linear matrix polynomial in u and the values and vectors of interest are its singular values and right kernel vectors. Finding these for an arbitrary matrix polynomial is discussed in the next section.

4 Hiding a Variable

An alternative way to obtain an overconstrained system from a well-constrained one is by *hiding* one of the original variables in the coefficient field. Hiding a variable instead of adding an extra polynomial possesses the advantage of keeping the number of polynomials constant. Our experience in solving polynomial systems in robotics and vision suggests that this usually leads to smaller eigenproblems. The resultant matrix is regarded as a matrix polynomial in the hidden variable and finding its singular values and kernel vectors generalizes the u -polynomial construction of the previous section.

Again we suppose the ideal is zero-dimensional and radical. Formally, given

$$f_0, \dots, f_n \in K[x_1, x_1^{-1}, \dots, x_{n+1}, x_{n+1}^{-1}] \quad (6)$$

we can view this system as

$$f_0, \dots, f_n \in (K[x_{n+1}])[x_1, x_1^{-1}, \dots, x_n, x_n^{-1}], \quad (7)$$

which is a system of $n+1$ Laurent polynomials in variables x_1, \dots, x_n . Notice that we have multiplied all polynomials by sufficiently high powers of x_{n+1} in order to

avoid dealing with denominators in the *hidden variable* x_{n+1} . This does not affect the system's roots in $(\overline{K}^*)^{n+1}$.

The sparse resultant of this system is a univariate polynomial in x_{n+1} . We show below how this formulation reduces the solution of the original well-constrained system to an eigenproblem.

Theorem 4.1 *Assume that M is a sparse resultant matrix for (7), with the polynomial coefficients specialized. Let $\alpha = (\alpha_1, \dots, \alpha_n) \in (\overline{K}^*)^n$ such that $(\alpha, \alpha_{n+1}) \in (\overline{K}^*)^{n+1}$ is a solution of $f_1 = \dots = f_{n+1} = 0$. Then $M(\alpha_{n+1})$ is singular and column vector $w = [\alpha^{q_1}, \dots, \alpha^{q_c}]$ lies in the right kernel of $M(\alpha_{n+1})$, where $\mathcal{E} = \{q_1, \dots, q_c\} \subset \mathbb{Z}^n$ are the exponent vectors indexing the columns of M .*

Proof For specialized polynomial coefficients, $M(\alpha_{n+1})$ is singular by definition. By construction, right multiplication by a vector of the column monomials specialized at a point produces a vector of the values of the row polynomials at this point. Let the i -th row of M contain the coefficients of $x^{p_i} f_j$, then

$$M(\alpha_{n+1})w = M(\alpha_{n+1}) \begin{bmatrix} \alpha^{q_1} \\ \vdots \\ \alpha^{q_c} \end{bmatrix} = \begin{bmatrix} \alpha^{p_1} f_{i_1}(\alpha, \alpha_{n+1}) \\ \vdots \\ \alpha^{p_c} f_{i_c}(\alpha, \alpha_{n+1}) \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}$$

where $i_1, \dots, i_c \in \{0, 1, \dots, n\}$.

□

Computationally it is preferable to have to deal with as small a matrix as possible. To this end we partition M into four blocks so that the upper left submatrix M_{11} is square, nonsingular and independent of x_{n+1} . Row and column permutations do not affect the matrix properties so we apply them to obtain a maximal M_{11} .

Gaussian elimination of the leftmost set of columns is now possible and expressed as matrix multiplication, where I is the identity matrix of appropriate size:

$$\begin{bmatrix} I & 0 \\ -M_{21}(x_{n+1})M_{11}^{-1} & I \end{bmatrix} \begin{bmatrix} M_{11} & M_{12}(x_{n+1}) \\ M_{21}(x_{n+1}) & M_{22}(x_{n+1}) \end{bmatrix} = \begin{bmatrix} M_{11} & M_{12}(x_{n+1}) \\ 0 & M'(x_{n+1}) \end{bmatrix}, \quad (8)$$

where

$$M'(x_{n+1}) = M_{22}(x_{n+1}) - M_{21}(x_{n+1})M_{11}^{-1}M_{12}(x_{n+1}).$$

Let $\mathcal{B} \subset \mathcal{E}$ index M' . We do not have *a priori* knowledge of the sizes of M_{11} and M' whether the subdivision or the incremental algorithm has constructed M .

Corollary 4.2 *Let $\alpha = (\alpha_1, \dots, \alpha_n) \in (\overline{K}^*)^n$ such that $(\alpha, \alpha_{n+1}) \in (\overline{K}^*)^{n+1}$ is a common zero of $f_0 = \dots = f_n = 0$. Then $\det M'(\alpha_{n+1}) = 0$ and, for any vector $v' = [\dots \alpha^q \dots]$, where q ranges over \mathcal{B} , $M'(\alpha_{n+1})v' = 0$.*

To recover the root coordinates, \mathcal{B} must affinely span \mathbb{Z}^n , otherwise we have to compute the kernel vector of matrix M which equals the concatenation of vectors v and v' , where v' is the kernel vector of M' and v is specified from (8):

$$\begin{aligned} \begin{bmatrix} M_{11} & M_{12}(\alpha_{n+1}) \\ 0 & M'(\alpha_{n+1}) \end{bmatrix} \begin{bmatrix} v \\ v' \end{bmatrix} &= \begin{bmatrix} 0 \\ 0 \end{bmatrix} \Rightarrow M_{11}v + M_{12}(\alpha_{n+1})v' = 0 \\ &\Leftrightarrow v = -M_{11}^{-1}M_{12}(\alpha_{n+1})v', \end{aligned}$$

since M_{11} is defined to be the maximal nonsingular submatrix. The concatenation $[v, v']$ is indexed by \mathcal{E} which always includes an affinely independent subset; because otherwise all n -fold mixed volumes are zero and no roots exist. Then, we recover all root coordinates by taking ratios of the vector entries.

We now concentrate on matrix polynomials and their companion matrices; for definitions and basic results consult [GLR82]. Denote the hidden variable by x , then we write

$$f_i \in (K[x])[x_1, x_1^{-1}, \dots, x_n, x_n^{-1}], \quad i = 0, \dots, n, \quad (9)$$

and we denote the univariate matrix $M'(x_{n+1})$ by $A(x)$. Let r be the size of A , and $d \geq 1$ the highest degree of x in any entry. We wish to find all values for x at which matrix

$$A(x) = x^d A_d + x^{d-1} A_{d-1} + \dots + x A_1 + A_0$$

becomes singular, where matrices A_d, \dots, A_0 are all square, of order r and have scalar entries. We refer to $A(x)$ as a *matrix polynomial* with degree d and matrix coefficients A_i . The values of x that make $A(x)$ singular are its *eigenvalues*. For every eigenvalue λ , there is a basis of the kernel of $A(\lambda)$ defined by the *right eigenvectors* of the matrix polynomial associated to λ . This is the eigenproblem for matrix polynomials, a classic problem in linear algebra.

If A_d is nonsingular then the eigenvalues and right eigenvectors of $A(x)$ are the eigenvalues and right eigenvectors of a *monic* matrix polynomial. Notice that this is always the case with the u -resultant formulation in the previous section.

$$A_d^{-1}A(x) = x^d I + x^{d-1} A_d^{-1} A_{d-1} + \dots + x A_d^{-1} A_1 + A_d^{-1} A_0,$$

where I is the $m \times m$ identity matrix. The *companion matrix* of this monic matrix polynomial is defined to be a square matrix C of order rd .

$$C = \begin{bmatrix} 0 & I & \cdots & 0 \\ \vdots & & \ddots & \\ 0 & 0 & \cdots & I \\ -A_d^{-1}A_0 & -A_d^{-1}A_1 & \cdots & -A_d^{-1}A_{d-1} \end{bmatrix}.$$

It is known that the eigenvalues of C are precisely the eigenvalues of the monic polynomial, whereas its right eigenvectors contain as subvectors the right eigenvectors of $A_d^{-1}A(x)$. Formally, assume $w = [v_1, \dots, v_d] \in \overline{K}^{md}$ is a (nonzero) right eigenvector of C , where each $v_i \in \overline{K}^m$, $i = 1, \dots, d$. Then v_1 is a (nonzero) right eigenvector of $A_d^{-1}A(x)$ and $v_i = \lambda^{i-1}v_1$, for $i = 2, \dots, d$, where λ is the eigenvalue of C corresponding to w .

We now address the question of a singular leading matrix in a non-monic polynomial. The following transformation is also used in the implementation in order to improve the conditioning of the leading matrix.

Lemma 4.3 *Assume matrix polynomial $A(x)$ is not identically singular for all x and let d be the highest degree in x of any entry. Then there exists a transformation $x \mapsto (t_1y + t_2)/(t_3y + t_4)$ for some $t_1, t_2, t_3, t_4 \in \mathbb{Z}$, that produces a new matrix polynomial $B(y)$ of the same degree and with matrices of the same dimension, such that $B(y)$ has a nonsingular leading coefficient matrix.*

It is easy to see that the resulting polynomial $B(y)$ has matrix coefficients of the same rank, for sufficiently generic scalars t_1, t_2, t_3, t_4 , since every matrix is the sum of $d + 1$ scalar products of A_i . Thus this transformation is often referred to as *rank balancing*.

Theorem 4.4 *If the values of the hidden variable in the solutions of $f_1 = \dots = f_{n+1} = 0$, which correspond to eigenvalues of the matrix polynomial, are associated to eigenspaces of unit dimension and $A(x)$ is not identically singular then we can reduce root-finding to an eigenproblem and some evaluations of the input polynomials at candidate roots.*

Proof We have seen that the new matrix polynomial $B(y)$ has a nonsingular leading coefficient. Moreover, finding its eigenvalues and eigenvectors is reduced to an eigenproblem of the companion matrix C . By hypothesis, the eigenspaces of C are

one-dimensional therefore for every root there is an eigenvector that yields n coordinates of the root by theorem 3.1. The associated eigenvalue may be either simple or multiple and yields the value of the hidden variable at the same root. Notice that extraneous eigenvectors and eigenvalues may have to be rejected by direct evaluation of the input polynomials at the candidate roots and a zero test. The right eigenvectors of $B(y)$ are identical to those of $A(x)$ but any eigenvalue λ of the former yields $(t_1\lambda + t_2)/(t_3\lambda + t_4)$ as an eigenvalue of $A(x)$. \square

The condition that all eigenspaces are unit-dimensional is equivalent to the solution coordinate at the hidden variable having unit geometric multiplicity. For this it suffices that the algebraic multiplicity of these solutions be one i.e., all hidden coordinates must be distinct. Since there is no restriction in picking which variable to hide, it is enough that one out of the original $n + 1$ variables have unit geometric multiplicity. If none can be found, we can specialize the hidden variable to each of the eigenvalues and solve every one of the resulting subsystems.

The asymptotic complexity bound for this method is given by the product of the bound in theorem 3.2 by d , the maximum polynomial degree in a single variable.

5 Implementation and Numerical Accuracy

Our implementation is entirely in Ansi C. The overall method has two stages, one online and one offline. A program of independent interest, which makes part of this package, is already available. It is the implementation of a fast algorithm for computing mixed volumes [EC95] and can be obtained by anonymous ftp on `robotics.eecs.Berkeley.edu` directory `MixedVolume`, or on `http://www.inria.fr/safir/emiris/mixvol`.

One advantage of the resultant method over previous algebraic as well as numerical methods is that the resultant matrix need only be computed once for all systems with the same set of exponents. So this step can often be done offline, while the eigenvalue calculations to solve the system for each coefficient specialization are online.

Another offline operation is to permute rows and columns of the given resultant matrix in order to create a maximal square submatrix at the upper left corner which is independent of the hidden variable. In order to minimize the computation involving polynomial entries and to reduce the size of the upper left submatrix that must be inverted, the program concentrates all constant columns to the left and within these columns permutes all zero rows to the bottom.

To find the eigenvector entries that will allow us to recover the root coordinates it is typically sufficient to examine \mathcal{B} indexing M_{22} and search for pairs of entries corresponding to exponent vectors v_1, v_2 such that $v_1 - v_2 = (0, \dots, 0, 1, 0, \dots, 0)$. This will let us compute the i -th coordinate if the unit appears at the i -th position. For inputs where we need to use points in $\mathcal{E} \setminus \mathcal{B}$, only the entries indexed by these points must be computed in vector v_α , in the notation of theorem 3.1. In any case the complexity of this step is dominated.

The input to the online solver is the set of all supports, the associated coefficients and matrix M whose rows and columns are indexed by monomial sets. The output is a list of candidate roots and the values of each at the given polynomials. The test for singularity of the matrix polynomial is implemented by substituting a few random values in the hidden variable and testing whether the resulting numeric matrix has full rank. For rational coefficients this is done in modular arithmetic so the answer is exact within a finite field, therefore nonsingularity is decided with very high probability.

Most of the online computation is numeric for reasons of speed; in particular we use double precision floating point numbers. We use the LAPACK package [ABB⁺92] because it implements state-of-the-art algorithms, including block algorithms on appropriate architectures, and provides efficient ways for computing condition numbers and error bounds. Moreover it is publicly available, portable, includes a C version and is soon to include a parallel version for shared-memory machines. Of course it is always possible to use other packages such as EISPACK [SBD⁺76] or LINPACK [BDMS79].

A crucial question in numerical computation is to predict the extent of roundoff error; see for instance [GL89]. To measure this we make use of the standard definition of matrix norm $\|A\|_p$ for matrix A . We define four condition numbers for A :

$$\kappa_p(A) = \|A\|_p \|A^{-1}\|_p, \quad p = 1, 2, \infty, F, \quad A \text{ square}, \quad \text{and if } A \text{ singular } \kappa_p(A) = \infty,$$

where F denotes the Frobenius norm. $\kappa_2(A)$ equals the ratio of the maximum over the minimum singular value of A . For a given matrix, the ratio of any two condition numbers with $p = 1, 2, \infty, F$ is bounded above and below by the square of the matrix dimension or its inverse. We also use $\|v\|_p$ for the p -norm of vector v .

As precise condition numbers are sometimes expensive to compute, LAPACK provides approximations of them and the associated error bounds. Computing these estimates is very efficient compared to the principal computation. These approximations run the risk of underestimating the correct values, though this happens extremely seldom in practice. Error bounds are typically a function of the condition

number and the matrix dimension; a reasonable assumption for the dependence on the latter is to use a linear function. Small condition numbers indicate well-behaved or *well-conditioned* matrices, e.g. orthogonal matrices are perfectly conditioned with $\kappa = 1$. Large condition numbers characterize *ill-conditioned* matrices.

After permuting rows and columns so that the maximal upper left submatrix is independent of u or the hidden variable x , we apply an LU decomposition with column pivoting to the upper left submatrix. We wish to decompose the maximal possible submatrix so that it has a reasonable condition number. The maximum magnitude of an acceptable condition number can be controlled by the user; dynamically, the decomposition stops when the pivot takes a value smaller than some threshold.

To compute M' we do not explicitly form M_{11}^{-1} but use its decomposition to solve linear problem $M_{11}X = M_{12}$ and then compute $M' = M_{22} - M_{21}X$. Different routines are used, depending on $\kappa(M_{11})$, to solve the linear problem, namely the slower but more accurate `dgesvx` function is called when this condition number is beyond some threshold. Let \hat{x}_j denote some column of X and let x_j be the respective column if no roundoff error were present. Then the error is bounded [GL89] by

$$\frac{\|x - \hat{x}\|_{\infty}}{\|x\|_{\infty}} \leq 4 \cdot 10^{-15} (|\mathcal{E}| - r) \kappa_{\infty}(M_{11}),$$

where $|\mathcal{E}| - r$ is the size of M_{11} and we have used $2 \cdot 10^{-16}$ as the machine precision under the IEEE double precision floating point arithmetic.

For nonsingular matrix polynomials $A(x)$ we try a few random integer quadruples (t_1, t_2, t_3, t_4) . We then redefine the matrix polynomial to be the one with lowest $\kappa(A_d)$. This operation of *rank balancing* is indispensable when the leading coefficient is nonsingular as well as ill-conditioned. Empirically we have observed that for matrices of dimension larger than 200, at least two or three quadruples should be tried since a lower condition number by two or three orders of magnitude is sometimes achieved. The asymptotic as well as practical complexity of this stage is dominated by the other stages.

If we manage to find a matrix polynomial with well-conditioned A_d we compute the equivalent monic polynomial and call the standard eigendecomposition routine. There are again two choices in LAPACK for solving this problem with an iterative or a direct algorithm, respectively implemented in routines `hsein` and `trevc`. Experimental evidence points to the former as being faster on problems where the matrix size is at least 10 times larger than the mixed volume, since an iterative solver can better exploit the fact that we are only interested in real eigenvalues and eigenvectors.

If A_d is ill-conditioned for all linear t_i transformations we build the matrix pencil and call the *generalized eigendecomposition* routine `dgegv` to solve $C_1x + C_0$. The latter returns pairs (α, β) such that matrix $C_1\alpha + C_0\beta$ is singular. For every (α, β) pair there is a nonzero right generalized eigenvector. For nonzero β we obtain the eigenvalues as α/β , while for zero β and nonzero α the eigenvalue tends to infinity and depending on the problem at hand we may or may not wish to discard it. The case $\alpha = \beta = 0$ occurs if and only if the pencil is identically zero within machine precision.

In recovering the eigenvector of matrix polynomial $A(x)$ from the eigenvector of its companion matrix, we can use any subvector of the latter. We choose the topmost subvector when the eigenvalue is smaller than the unit, otherwise we use one of the lower subvectors. Nothing changes in the algorithm, since ratios of the entries will still yield the root, yet the stability of these ratios is improved.

There are certain properties of the problem that have not been exploited yet. Typically, we are interested only in real solutions. We could concentrate, therefore, on the real eigenvalues and eigenvectors and choose the algorithms that can distinguish between them and complex solutions at the earliest stage. Moreover, bounds on the root magnitude may lead to substantial savings, as exemplified in [Man94]. In the rest of this article we examine these issues in the light of concrete applications of our program.

6 Camera Motion from Point Matches

Camera motion reconstruction, or *relative orientation*, in its various forms is a basic problem in photogrammetry, including the computation of the shape of an object from its motion. Formally, we are interested in the problem of computing the displacement of a camera between two positions in a static environment. Given are the coordinates of certain points in the two views under perspective projection on calibrated cameras. Equivalently, this problem consists in computing the displacement of a rigid body, whose identifiable features include only points, between two snapshots taken by a stationary camera. We consider, in particular, the case where the minimum number of 5 point matches is available. In this case the algebraic problem reduces to a well-constrained system of polynomial equations and we are able to give a closed-form solution.

Typically, computer vision applications use at least 8 points in order to reduce the number of possible solutions to 3 and, for generic configurations, to one. In addition, computing the displacement reduces to a linear problem and the effects

system	operation	CPU time
6×6	mixed volume	1m 16s
6×5	sparse resultant (offline)	12s
6×6 (first)	root finding (online)	0.2s
6×6 (second)	root finding (online)	1s (SUN SPARC 20)

Table 1: Camera motion from point matches: running times are measured on a DEC ALPHA 3000 except for the second system which is solved on a SUN SPARC 20.

of noise in the input can be diminished [LH81]. Our approach shows performance comparable to these methods and, as it requires the minimum number of data points, it is well-suited for detecting and eliminating *outliers* in the presence of noise. These are data points which are so much affected by noise that they should not be taken into account.

6.1 Algebraic Formulation

To formalize, let orthonormal 3×3 matrix $R \in \text{SO}(3, \mathbb{R})$ denote the rotation. Let (column) vector $t \in \mathbb{R}^3$ denote the camera translation in the original frame of reference. The 5 points in the two images are (column) vectors $a_i, a'_i \in \mathbb{P}_{\mathbb{R}}^2$, for $i = 1, \dots, 5$ in the first and second frame of reference respectively. It is clear that the magnitude of t cannot be recovered, hence there are 5 unknowns, 3 defining the rotation and 2 for the translation.

The following quaternion formulation was independently suggested by J. Canny and [Hor91]. Let x, y, z be 3-vectors and $\dot{x} = [x_0, x], \dot{y} = [y_0, y], \dot{z} = [z_0, z] \in \mathbb{R}^4$ be arbitrary quaternions. Let $\dot{x}\dot{y}$ represent a quaternion product and $\dot{z}^* = [z_0, -z]$ be the conjugate quaternion of \dot{z} . Quaternions $\dot{q} = [q_0, q], \dot{t} = [t_0, t] \in \mathbb{R}^4$ represent the rotation and translation respectively. A rotation represented by angle ϕ and unit 3-vector s is expressed uniquely by quaternion $[\cos \phi/2, \sin \phi/2 s]$, hence any rotation quaternion has unit 2-norm.

The equations in terms of the quaternions are homogeneous. After dehomogenization we obtain 6 polynomials in 6 variables organized in two 3-vectors. We denote these vectors by q, d . If $i = 1, \dots, 5$, the six equations are:

$$\begin{aligned} (a_i^T q)(d^T a'_i) + a_i^T a'_i + (a_i \times q)^T a'_i + (a_i \times q)^T (d \times a'_i) + a_i^T (d \times a'_i) &= 0, \\ 1 - d^T q &= 0 \quad (10) \end{aligned}$$

machine	clock rate	memory	SpecInt92	SpecFP92
DEC ALPHA 3000/300	150 MHz	64 MB	67	77
SUN SPARC 20/61	60 MHz	32 MB	95	93

Table 2: Hardware specifications

This system is well-constrained so we can apply Bernstein’s bound to approximate the number of roots in $(\mathbb{C}^*)^6$. The mixed volume of this system is 20, which is an exact bound [Dem88]. The performance of our implementation on mixed volume is shown in table 1 for the DEC ALPHA 3000 of table 2.

6.2 Applying the Resultant Solver

Resultant matrix M is of dimension 60, while only 20 columns contain hidden variable q_1 . 40×40 submatrix M_{11} is inverted and the resulting 20×20 pencil is passed to an eigendecomposition routine. The running times for the two main phases are reported in table 1. The rest of this section examines the accuracy of our procedure on two specific instances from [FM90]. We use error criterion

$$\sum_{i=1}^5 \frac{1}{\|a_i\| \cdot \|a'_i\|} |a_i^T (t \times R a'_i)|, \quad (11)$$

where $|\cdot|$ denotes absolute value. $\|\cdot\|$ is the vector 2-norm and t, R are the calculated translation and rotation. This expression vanishes when the solutions to R, t are exact, otherwise it returns the absolute value of error normalized by the input norm.

The first example admits the maximum of 10 pairs of real solutions. M_{11} is well-conditioned with $\kappa < 10^3$, so it is factorized to yield an optimal pencil of dimension 20. The latter has a well-conditioned leading coefficient with $\kappa < 10^4$, so it yields a monic polynomial on which the standard eigendecomposition is applied. The maximum value of error criterion (11) is less than $1.3 \cdot 10^{-6}$, which is very satisfactory.

The input parameters are sufficiently generic in order to lead to the maximum number of real solutions. This explains the good conditioning of the matrices. The second example looks at a less generic instance: namely a configuration that has only 8 distinct real solutions which are clustered close together. This is manifest in the fact that matrix M_{11} is ill-conditioned. The input is constructed by applying a rotation of 60.00145558° about the z axis in the first frame and a translation of $(0.01, 0.01, -1)$.

Exactly the same procedure is used as before to produce a 60×60 resultant matrix, but now the upper leftmost 40×40 submatrix M_{11} has $\kappa > 0.6 \cdot 10^5$ so we choose not to invert it. Instead, the 60×60 pencil is considered: the leading matrix has $\kappa > 10^8$ in its original form and for any of 4 random transformations, hence it is not inverted and the generalized eigendecomposition is applied.

Applying error criterion (11), the largest absolute value is $7.4 \cdot 10^{-5}$, hence all 16 real solutions are quite accurate. There are another 4 complex solutions and 40 infinite ones. The total CPU running time for the online phase is, on the average, 1 second on the SUN SPARC 20 of table 2.

It is interesting to observe in connection to roots with zero coordinates, that in the second example our solver recovers roots in $\mathbb{C}^n \setminus (\mathbb{C}^*)^n$, namely we find a camera motion whose rotation quaternion \hat{q} has $q_1 = q_2 = 0$. Such roots can be thought of as limits of roots in $(\mathbb{C}^*)^n$ as the system coefficients deform. As long as the variety does not generically reside in $\mathbb{C}^n \setminus (\mathbb{C}^*)^n$, roots with zero coordinates will always be recovered. This is typically the case in practical applications. For the particular example, there is a stronger reason why all roots are recovered, namely all polynomials include a constant term (see theorem 2.5).

There exist various implementations of linear methods requiring at least 8 points, including Luong's [Luo92]. We have been able to experiment with this program which implements the least-squares method of [TH84], and found it faster on both instances but less accurate on the second one. In particular, we chose specific solutions to generate an additional 3 matches for each of the problems above. The average CPU time on the two examples is 0.08 seconds on a SUN SPARC 20. On the first example the output is accurate to at least 7 digits. On the second example, Luong's implementation returns a rotation of 60.0013° about $(-10^{-5}, 10^{-5}, 1)$ and a unit translation vector of $(-10^{-5}, -10^{-5}, -1)$. But the latter differs significantly from the true vector $t = (0.01, 0.01, -1)$.

7 Conformational Analysis of Cyclic Molecules

A relatively new branch of computational biology has been emerging as an effort to apply successful paradigms and techniques from geometry and robot kinematics to predicting the structure of molecules, embedding them in euclidean space and finding the energetically favorable configurations [PC94, Emi94, EM96]. The main premise for this interaction is the observation that various structural requirements on molecules can be modeled as geometric or kinematic constraints.

This section examines the problem of computing all *conformations* of a *cyclic* molecule, which reduces to an *inverse kinematics* problem. Conformations specify the 3-dimensional structure of the molecule. It has been argued by Gō and Scherga [GS70] that energy minima can be approximated by allowing only the dihedral angles to vary, while keeping bond lengths and bond angles fixed. At a first level of approximation, therefore, solving for the dihedral angles under the assumption of *rigid geometry* provides information for the energetically favorable configurations.

We consider molecules of six atoms to illustrate our approach and show that the corresponding algebraic formulation conforms to our model of sparseness. Our resultant solver is able to compute all solutions accurately even in cases where multiple solutions exist. An in-depth study of this problem, including a comparison of different resultant-based methods for system solving, has been recently presented in [EM96].

7.1 Algebraic Formulation

The molecule has a cyclic backbone of 6 atoms, typically of carbon. They determine primary structure, the object of our study. Carbon-hydrogen or other bonds outside the backbone are ignored. The bond lengths and angles provide the constraints while the six dihedral angles are allowed to vary. In kinematic terms, atoms and bonds are analogous to *links* and *joints* of a *serial mechanism* in which each pair of consecutive axes intersects at a link. This implies that the link offsets are zero for all six links which allows us to reduce the 6-dimensional problem to a system of 3 polynomials in 3 unknowns. The product of all link transformation matrices is the identity matrix, since the end-effector is at the same position and orientation as the base link.

We adopt an approach proposed by D. Parsons [Par94]. Notation is defined in figure 1. Backbone atoms are regarded as points $p_1, \dots, p_6 \in \mathbb{R}^3$; the unknown dihedrals are the angles $\omega_1, \dots, \omega_6$ about axes (p_6, p_1) and (p_{i-1}, p_i) for $i = 2, \dots, 6$. For readers familiar with the kinematics terminology, the Denavit-Hartenberg parameters are

$$\alpha_i = 180^\circ - \phi_i, d_i = L_i, a_i = 0, \theta_i = \omega_i.$$

Each of triangles $T_1 = \Delta(p_1, p_2, p_6)$, $T_2 = \Delta(p_2, p_3, p_4)$ and $T_3 = \Delta(p_4, p_5, p_6)$ is fixed for constant bond lengths L_1, \dots, L_6 and bond angles ϕ_1, ϕ_3, ϕ_5 . Then the lengths of (p_2, p_6) , (p_2, p_4) and (p_4, p_6) are constant, hence *base* triangle $\Delta(p_2, p_4, p_6)$ is fixed in space, defining the *xy*-plane of a coordinate frame. Let θ_1 be the (dihedral) angle between the plane of $\Delta(p_1, p_2, p_6)$ and the *xy*-plane. Clearly, for any conformation

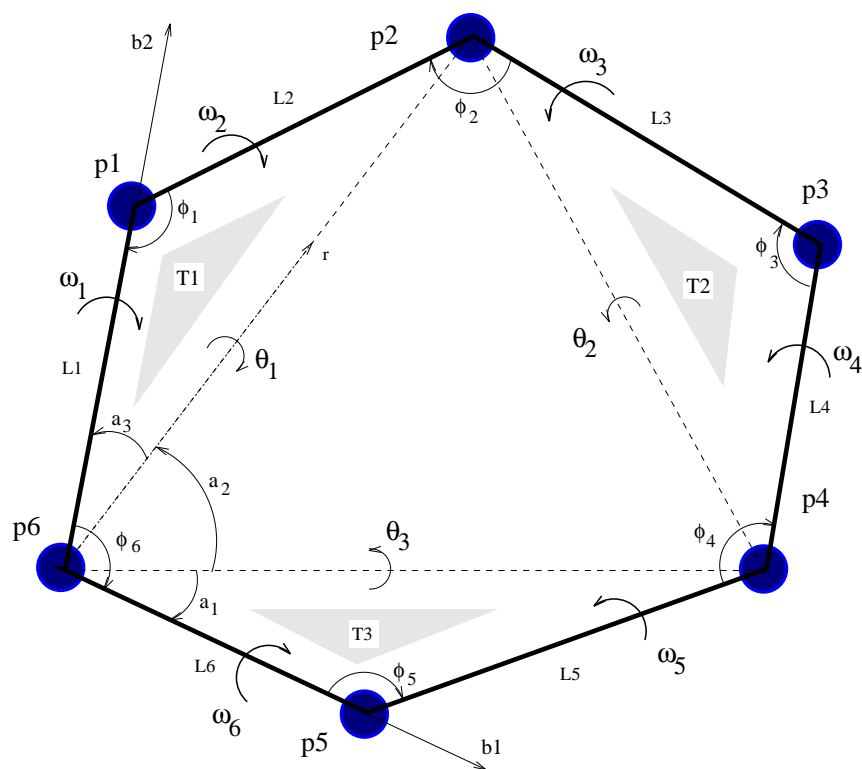


Figure 1: The cyclic molecule.

θ_1 is well-defined. Similarly we define angles θ_2 and θ_3 , as shown in figure 1. We call them *flap* (dihedral) angles to distinguish them from the bond dihedrals.

Conversely, given lengths L_i , angles ϕ_i for $i = 1, \dots, 6$ and flap angles θ_i for $i = 1, \dots, 3$ the coordinates of all points p_i are uniquely determined and hence the bond dihedral angles and the associated conformation are all well-defined. We have therefore reduced the problem to computing the three flap angles θ_i which satisfy the constraints on bond angles ϕ_2, ϕ_4, ϕ_6 .

Hence we obtain polynomial system

$$\begin{aligned}
\alpha_{11} + \alpha_{12} \cos \theta_2 + \alpha_{13} \cos \theta_3 + \alpha_{14} \cos \theta_2 \cos \theta_3 + \alpha_{15} \sin \theta_2 \sin \theta_3 &= 0, \\
\alpha_{21} + \alpha_{22} \cos \theta_3 + \alpha_{23} \cos \theta_1 + \alpha_{24} \cos \theta_3 \cos \theta_1 + \alpha_{25} \sin \theta_3 \sin \theta_1 &= 0, \\
\alpha_{31} + \alpha_{32} \cos \theta_1 + \alpha_{33} \cos \theta_2 + \alpha_{34} \cos \theta_1 \cos \theta_2 + \alpha_{35} \sin \theta_1 \sin \theta_2 &= 0, \\
\cos^2 \theta_1 + \sin^2 \theta_1 - 1 &= 0, \\
\cos^2 \theta_2 + \sin^2 \theta_2 - 1 &= 0, \\
\cos^2 \theta_3 + \sin^2 \theta_3 - 1 &= 0,
\end{aligned} \tag{12}$$

where the α_{ij} are input coefficients.

This system has Bézout bound of 64 and mixed volume 16; the mixed volume is the exact number of complex roots generically as we shall prove below by demonstrating an instance with 16 real roots. Notice that 16 is also the exact number of solutions, generically, to the *general* inverse kinematics problem with 6 rotational joints (6R).

For our resultant solver we prefer an equivalent formulation with a smaller number of polynomials, obtained by applying the standard transformation to half-angles that gives *rational* equations in the new unknowns t_i :

$$t_i = \tan \frac{\theta_i}{2} : \quad \cos \theta_i = \frac{1 - t_i^2}{1 + t_i^2}, \quad \sin \theta_i = \frac{2t_i}{1 + t_i^2}, \quad i = 1, 2, 3.$$

This transformation captures automatically the last three equations in (12). By multiplying both sides of the i -th equation by $(1 + t_j^2)(1 + t_k^2)$, where (i, j, k) is a permutation in $S(1, 2, 3)$, the polynomial system becomes

$$\begin{aligned}
f_1 &= \beta_{11} + \beta_{12}t_2^2 + \beta_{13}t_3^2 + \beta_{14}t_2^2t_3^2 + \beta_{15}t_2t_3 = 0 \\
f_2 &= \beta_{21} + \beta_{22}t_3^2 + \beta_{23}t_1^2 + \beta_{24}t_3^2t_1^2 + \beta_{25}t_3t_1 = 0 \\
f_3 &= \beta_{31} + \beta_{32}t_1^2 + \beta_{33}t_2^2 + \beta_{34}t_1^2t_2^2 + \beta_{35}t_1t_2 = 0
\end{aligned} \tag{13}$$

where β_{ij} are input coefficients. The new system has again Bézout bound of 64 and mixed volume 16.

7.2 Applying the Resultant Solver

The first instance is a synthetic example for which we fix one feasible conformation with all flap angles equal to 90° . All polynomials are multiplied by 8 in order for the coefficients to be all integers, then β_{ij} is the (i, j) -th entry of matrix

$$\begin{bmatrix} -9 & -1 & -1 & 3 & 8 \\ -9 & -1 & -1 & 3 & 8 \\ -9 & -1 & -1 & 3 & 8 \end{bmatrix}.$$

The symmetry of the problem is bound to produce root coordinates of high multiplicity, so we decide to follow the first approach to solving the system (sect. 3) and add polynomial

$$f_0 = u + 31t_1 - 41t_2 + 61t_3$$

with randomly selected coefficients. In this system, the 3-fold mixed volumes are 12, 12, 12, 16 hence the sparse resultant has total degree 52 and degree 16 in f_0 . The resultant matrix is regular and has dimension 86, with 30 rows corresponding to f_0 . This is the offline phase; the rest corresponds to the online execution of the solver.

The entire 56×56 upper left submatrix is decomposed and is relatively well-conditioned. In the 30×30 matrix polynomial, the leading matrix coefficient is singular within machine precision; two random transformations are used but fail to improve significantly the conditioning of the matrix. Therefore the generalized eigenproblem routine is called on the 30×30 pencil and produces 12 complex solutions, 3 infinite real solutions and 15 finite real roots. The absolute value of the four polynomials on the candidate values lies in $[0.6 \cdot 10^{-9}, 0.3 \cdot 10^{-3}]$ for values that approximate true solutions and in $[7.0, 3.0 \cdot 10^{20}]$ for spurious answers. Our program computes the true roots to at least 5 digits as seen by comparing with the exact solutions computed by MAPLE V using Gröbner bases over the rationals. The latter are

$$\pm(1, 1, 1), \pm(5, -1, -1), \pm(-1, 5, -1), \pm(-1, -1, 5).$$

The average CPU time of the online phase on the SUN SPARC 20 of table 2 is 0.4 seconds.

Usually noise enters in the process that produces the coefficients; this example models this phenomenon. We consider the *cyclohexane* molecule which has 6 carbon atoms at equal distances and equal bond angles. Starting with the pure cyclohexane,

we randomly perturb them by about 10% to obtain β_{ij} as the entries of matrix

$$\begin{bmatrix} -310 & 959 & 774 & 1313 & 1389 \\ -365 & 755 & 917 & 1269 & 1451 \\ -413 & 837 & 838 & 1352 & 1655 \end{bmatrix}.$$

We used the second approach to define an overconstrained system, namely by hiding variable t_3 in the coefficient field (sect. 4). The resultant matrix has dimension 16 and is quadratic in t_3 , whereas the 2-fold mixed volumes are all 4 and the sparse resultant has degree $4 + 4 + 4 = 12$.

The monic quadratic polynomial reduces to a 32×32 companion matrix on which the standard eigendecomposition is applied. After rejecting false candidates, the recovered roots cause the maximum absolute value of the input polynomials to be 10^{-5} . We check the computed solutions against those obtained by a Gröbner bases computation over the integers and observe that each contains at least 8 correct digits. The total CPU time on a SUN SPARC 20 is 0.2 seconds on average for the online phase.

Lastly we report on an instance where the input parameters are sufficiently generic to produce 16 real roots. The β_{ij} coefficients are given by matrix

$$\begin{bmatrix} -13 & -1 & -1 & -1 & 24 \\ -13 & -1 & -1 & -1 & 24 \\ -13 & -1 & -1 & -1 & 24 \end{bmatrix}.$$

We hide t_3 and arrive at a resultant matrix of dimension 16, whereas the sparse resultant has degree 12. The monic polynomial and the companion matrix are of dimension 32. There are 16 real roots. Four of them correspond to eigenvalues of unit geometric multiplicity, while the rest form four groups, each corresponding to a triple eigenvalue. For the latter the eigenvectors give us no valid information, so we recover the values of t_1, t_2 by looking at the other solutions and by relying on symmetry arguments. The computed roots are correct to at least 7 decimal digits. The average CPU time of the online part is 0.2 seconds on a SUN SPARC 20.

8 Conclusion

We have examined several computational aspects of sparse elimination theory and, in particular, the use of sparse resultant matrices for reducing root-finding to an eigenproblem. A general solver has been implemented based on this approach and has

been applied successfully to fundamental problems in vision, robot kinematics and structural biology. These problems are of moderate size and exhibit sparse structure as modeled by the Newton polytopes and the mixed volume. The efficiency and accuracy of our solver imply that sparse elimination may be the method of choice for such systems.

Automating the different ways to deal with numerically unstable inputs will improve the implementation. For instance, clustering neighboring eigenvalues and computing the error on the average value significantly improves accuracy. A question of practical as well as theoretical interest is to handle the case of repeated roots efficiently.

Acknowledgment

I would like to thank John Canny, my thesis advisor, for many insightful suggestions. Most of this work was conducted as part of my Ph.D. thesis in the Computer Science Division of U.C. Berkeley. Financial support acknowledged through a Packard Foundation Fellowship and NSF PYI Grant IRI-8958577.

References

- [ABB⁺92] E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, S. Ostrouchov, and D. Sorensen. *LAPACK Users' Guide*. SIAM, Philadelphia, 1992.
- [BDMS79] J. Bunch, J. Dongarra, C. Moler, and G.W. Stewart. *LINPACK User's Guide*. SIAM, Philadelphia, 1979.
- [Ber75] D.N. Bernstein. The number of roots of a system of equations. *Funct. Anal. and Appl.*, 9(2):183–185, 1975.
- [CE93] J. Canny and I. Emiris. An efficient algorithm for the sparse mixed resultant. In G. Cohen, T. Mora, and O. Moreno, editors, *Proc. Intern. Symp. on Applied Algebra, Algebraic Algor. and Error-Corr. Codes, Lect. Notes in Comp. Science 263*, pages 89–104, Puerto Rico, 1993. Springer.
- [CR91] J. Canny and J.M. Rojas. An optimal condition for determining the exact number of roots of a polynomial system. In *Proc. ACM Intern. Symp. on Symbolic and Algebraic Computation*, pages 96–102, Bonn, July 1991.
- [CW90] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *J. Symbolic Computation*, 9:251–280, 1990.

- [Dem88] M. Demazure. Sur deux problèmes de reconstruction. Technical Report 882, I.N.R.I.A., Rocquencourt, 1988.
- [EC95] I.Z. Emiris and J.F. Canny. Efficient incremental algorithms for the sparse resultant and the mixed volume. *J. Symbolic Computation*, 20(2):117–149, August 1995.
- [EM96] I.Z. Emiris and B. Mourrain. Polynomial system solving: The case of a 6-atom molecule. Technical Report 3075, INRIA Sophia-Antipolis, France, December 1996.
- [Emi94] I.Z. Emiris. *Sparse Elimination and Applications in Kinematics*. PhD thesis, Computer Science Division, Univ. of California at Berkeley, December 1994.
- [Emi96] I.Z. Emiris. On the complexity of sparse elimination. *J. Complexity*, 12:134–166, 1996.
- [ER94] I.Z. Emiris and A. Rege. Monomial bases and polynomial system solving. In *Proc. ACM Intern. Symp. on Symbolic and Algebraic Computation*, pages 114–122, Oxford, July 1994.
- [FM90] O.D. Faugeras and S. Maybank. Motion from point matches: Multiplicity of solutions. *Intern. J. Comp. Vision*, 4:225–246, 1990.
- [Ful93] W. Fulton. *Introduction to Toric Varieties*. Number 131 in Annals of Mathematics. Princeton University Press, Princeton, 1993.
- [GL89] G.H. Golub and C.F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, Maryland, 1989.
- [GLR82] I. Gohberg, P. Lancaster, and L. Rodman. *Matrix Polynomials*. Academic Press, New York, 1982.
- [GS70] N. Gö and H.A. Scheraga. Ring closure and local conformational deformations of chain molecules. *Macromolecules*, 3(2):178–187, 1970.
- [Hor91] B.K.P. Horn. Relative Orientation Revisited. *J. Opt. Soc. Am.*, 8(10):1630–1638, 1991.
- [HS95] B. Huber and B. Sturmfels. A polyhedral method for solving sparse polynomial systems. *Math. Comp.*, 64(212):1542–1555, 1995.
- [HS96] B. Huber and B. Sturmfels. Bernstein’s theorem in affine space. *Discr. and Computational Geometry*, 1996. To appear.
- [LH81] H.C. Longuet-Higgins. A Computer Algorithm for Reconstructing a Scene from Two Projections. *Nature*, 293:133–135, 1981.
- [Luo92] Q.-T. Luong. *Matrice Fondamentale et Auto-calibration en Vision par Ordinateur*. PhD thesis, Université de Paris-Sud, Orsay, Dec. 1992.

-
- [Man94] D. Manocha. Computing selected solutions of polynomial equations. In *Proc. ACM Intern. Symp. on Symbolic and Algebraic Computation*, pages 1–8, 1994.
- [Par94] D. Parsons, 1994. Personal Communication.
- [PC94] D. Parsons and J. Canny. Geometric problems in molecular biology and robotics. In *Proc. 2nd Intern. Conf. on Intelligent Systems for Molecular Biology*, pages 322–330, Palo Alto, CA, August 1994.
- [PS93] P. Pedersen and B. Sturmfels. Product formulas for resultants and Chow forms. *Math. Zeitschrift*, 214:377–396, 1993.
- [SBD⁺76] B.T. Smith, J.M. Boyle, J.J. Dongarra, B.S. Garbow, Y. Ikebe, V.C. Klema, and C.B. Moler. *Matrix Eigensystem Routines – EISPACK Guide*. Lect. Notes in Comp. Science, 6. Springer, Berlin, 1976.
- [TH84] R.Y. Tsai and T.S. Huang. Uniqueness and estimation of three-dimensional motion parameters of rigid objects with curved surfaces. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 6:13–27, 1984.
- [VVC94] J. Verschelde, P. Verlinden, and R. Cools. Homotopies exploiting Newton polytopes for solving sparse polynomial systems. *SIAM J. Numerical Analysis*, 31(3):915–930, 1994.



Unité de recherche INRIA Lorraine, Technopôle de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS LÈS NANCY
Unité de recherche INRIA Rennes, Irista, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unité de recherche INRIA Rhône-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unité de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unité de recherche INRIA Sophia Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA ANTIPOLIS Cedex

Éditeur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
ISSN 0249-6399