Optimized parametrization of systems of incidences between rigid bodies

Meera Sitharam [1,2,3], Jörg Peters [1,4], Yong Zhou [2,5]

Abstract

Graphs of pairwise incidences between collections of rigid bodies occur in many practical applications and give rise to large algebraic systems for which all solutions have to be found. Such pairwise incidences have explicit, simple and rational parametrizations that, in principle, allow us to partially resolve these systems and arrive at a reduced, parametrized system in terms of the rational parameters. However, the choice of incidences and the partial order of incidence resolution strongly determine the algebraic complexity of the reduced, parametrized system – measured primarily in the number of variables and secondarily in the degree of the equations.

Using a pairwise *overlap graph*, we introduce a combinatorial *class* of *incidence tree* parametrizations for a collection of rigid bodies. Minimizing the algebraic complexity over this class reduces to a purely combinatorial optimization problem that is a special case of the *set cover* problem. We quantify the exact improvement of algebraic complexity obtained by optimization and illustrate the improvement by examples that can not be solved without optimization.

Since incidence trees represent only a subclass of possible parametrizations, we characterize when optimizing over this class is useful. That is, we show what properties of *standard collections of rigid bodies* are necessary for an optimal incidence tree to have minimal algebraic complexity. For a standard collections of rigid bodies, the optimal incidence tree parameterization offers lower algebraic complexity than any other known parameterization.

Key words: Partial non-linear elimination, Optimal parametrization, Constraint graphs, Decomposition-Recombination of geometric constraint systems, Combinatorial rigidity, Systems of Linkages

[1] University of Florida
[2] Work supported in part by NSF Grant CCF 04-0404116, NSF Grant EIA 02-18435
[3] corresponding author: sitharam@cise.ufl.edu
[4] Work supported in part by NSF Grant DMI-0400214, CCF-0430891
[5] Amazon Corporation

1 Introduction and Motivation

A well-known approach in molecular conformation [4] and in kinematics is to model a *chain or cycle of pairwise incidences* representing molecular bonds or articulated robotic links in terms of 'half-angle formulas' (see e.g. [19,14]). Each formula fixes the link except for a rotation about one axis: It is a rational parameterization, based on the stereographic projection, that encodes the relative degrees of freedom between two rigid bodies that share an axis defined by two points. If the two rigid bodies share only one point, the resulting three degrees of freedom can be rationally parameterized via a less well-known quaternion motion. Completing the picture, if the two rigid bodies share three points, they generically form one of two possible rigid bodies.

Challenge and Scope. Our goal is to resolve pairwise incidences between rigid bodies in a more general setting than chains or cycles, namely for whole *graphs of rigid body interactions* (see e.g. Figure 1). Such graphs are prevalent, for example, in the modeling of protein backbones [21,10]: protein data bank (pdb) files now contain information about how backbones decompose into rigid parts. Such graphs also occur in constraint systems from mechanical CAD and kinematics (see e.g. the survey paper [15]). We are interested in the case where the entire system or collection of rigid bodies generically has at most a *finite* number of distinct *real* solutions. In this case, the entire system is called *generically rigid* [7] and if real solutions exist, we obtain several distinct instances of the corresponding composite rigid body. In the above applications, especially in molecular modeling, we are interested in finding all of these solutions. If the rigid system generically has at least one solution, it is called generically *well-constrained* (minimally rigid, in the terminology of combinatorial rigidity).

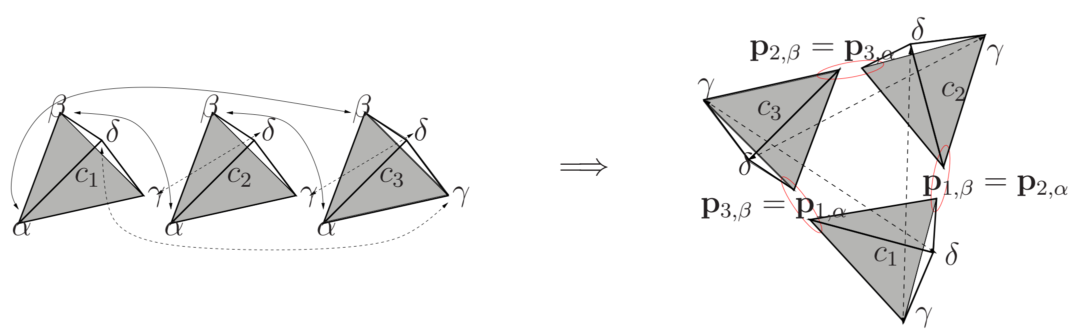

Fig. 1. Problem `tetra`. (*left*) Three rigid bodies (tetrahedra) c_i, $i = 1, 2, 3$ with points $\mathbf{p}_{i,j} \in \mathbb{R}^3$ for $j \in \{\alpha, \beta, \gamma, \delta\}$. *Solid* arrow-curves indicate incidences, i.e., the tetrahedra are to be moved so that corresponding points coincide. *Dashed* curves indicate fixed length bars between pairs of these points. These define three further rigid bodies c_4, c_5, c_6 (cf. Figure 2). (*right*) shows a possible assembly of the tetrahedra for a specific initialization and choice of fixed lengths within the rigid bodies.

Intuitively, for any collection of rigid bodies that interact via a graph of pairwise incidences, a *tree* of these pairs can be resolved, one pair at a time, by attaching to and representing one rigid body in the other's coordinates using rational parameterizations. The resolution of such a pairwise incidence tree yields a partial ordering of elimination steps and generates a parametrized system of remaining constraints expressed in terms of the remaining degrees of freedom. The remaining degrees of freedom are the *parameters* of the rational parameterizations used to resolve the tree. The resulting *parametrized system* stands a much better chance of being solved by algebraic-numeric solvers. Note that the resulting parametrized system depends on the tree, its traversal and the (incidence tree) parameterizations employed.

Contribution and Organization. The key contributions of this paper are

- The introduction of the concept of *incidence tree parameterizations* for resolving collections of rigid bodies.
- Conditions that define a *standard collection of rigid bodies* so that the optimal incidence tree parameterization optimizes algebraic complexity over all known parameterizations; and an effective strategy for standardizing collections.
- Three detailed illustrations.

Section 2 motivates the paper by example: it reviews the unoptimized approach to solving a well-constrained system of pairwise incidences among a collection S of rigid bodies; the reduction in complexity by applying well-chosen (pair-wise incidence) parameterizations; and finally, the further dramatic reduction by considering the *complete* collection of (maximal) rigid bodies and properly choosing and partially ordering its pairwise incidences. This example motivates an optimization problem over a natural class of rationally parametrized polynomial systems, each of which represents the same well-constrained system S. This class of *incidence tree parametrizations* is formally defined in Section 4.

While Section 2 shows that incidence tree parameterizations can be highly effective, Section 3 makes the important point that incidence tree parametrizations can be inefficient if we do not first *standardize* the collection. Conversely, once the collection is standard, the optimal incidence tree parameterization offers lower algebraic complexity than any other known parameterization. Characterizing standard collections of rigid bodies and establishing their importance is a key contribution of the paper and yields an effective strategy for dealing with non-standard collections: to alternate the standardization of the collection with the optimization of Section 4.1, in recursive stages. Fortunately, as pointed out in Section 4.5, such recursive standardization is part of the decomposition phase of efficient decomposition-recombination algorithms (DR-planners) for general geometric constraint systems [15,18]. In typical applications (CAD and molecular modeling) where collections of rigid body incidences occur, DR-planners are commonly used so that collections can be assumed to be standard.

In Section 4, Definition 4, we define a pairwise *overlap graph* to represent standard collections of rigid bodies. We use it to easily compute the algebraic complexity (number of variables and degree) of the collection's incidence tree parametrizations. Minimizing the algebraic complexity over this class of incidence trees for standard collections of rigid bodies is then a purely combinatorial problem, a tractable, special case of the *set cover* problem. The special case can be solved with a straightforward algorithm (Section 4.1). We quantify the improvement in algebraic complexity due to minimization over the unoptimized system (Section 4.2); and the effect on the overall computational complexity of resolving the collection of rigid bodies (Section 4.3).

We apply the developed theory to three natural examples of standard collections of rigid bodies (Section 5). The examples
(a) illustrate the improvement in algebraic complexity of the optimized incidence tree parametrized system (the examples could not be solved by current algebraic-numeric without optimization); and
(b) suggests (Section 6) that this parametrization (i) illuminates the solution space structure for incident standard collections of rigid bodies; (ii) could be used for optimizing the algebraic complexity of more general algebraic systems.

2 Parameterizing and Resolving Collections of Rigid Bodies

In this section, we illustrate by example some parameterizations of pairwise incidences that considerably reduce the complexity of resolving a well-constrained collection of incident rigid bodies into a single rigid body (Section 2.2). And we illustrate the additional reduction of complexity resulting from optimizing the partial order of elimination and parametrization (Section 2.3).

To motivate the framework, and show the impact of parameterization and then optimization we first describe an unoptimized, unparameterized formulation (Section 2.1) and explain the incidence elimination.

Example 1 (Problem `tetra`) *Figure 1 (left) shows a* collection $c := \{c_i\}$ *of (rigid) tetrahedra* c_1, c_2, c_3 *and (rigid) bars* c_4, c_5, c_6*, constrained by a system* $\mathcal{I}$ *of* incidences*, corresponding to shared points (see* (1) *below). The actual positions* $\mathbf{p}_{i,j}$ *and distance values* d_i *within the rigid bodies are not relevant for the discussion in this section. They are simply constants in the polynomial equations (that allow for a solution). The collection is* generically well-constrained *in 3D. Section 5 will present specific choices and numerical solutions.*

Obtaining a *realization* or resolution of c means fixing a *home coordinate system* h, say that of c_1, and repositioning $c_2, \ldots, c_6$ in the coordinate system of h in such a way that the incidences are satisfied. We denote by

$$\mathbf{p}_{i,j} \in \mathbb{R}^3 \quad \text{the position of the } j\text{th point in the } i\text{th rigid body}, c_i.$$

The challenge can be reduced, for example, to positioning the 2 remaining tetrahedra, while satisfying the incidences, and treating the fixed-length bars as distance constraints. I.e., for each of the 2 remaining tetrahedra we want to *find* a translation $T_i \in \mathbb{R}^3$ and the parameters of a matrix $M_i \in \mathbb{R}^{3\times 3}$, representing the composition of three rotations, so that for each given point j in tetrahedron c_i coinciding with point j' in tetrahedron $c_{i'}$

$$M_i\ \mathbf{p}_{i,j} + T_i = M_{i'}\ \mathbf{p}_{i',j'} + T_{i'} \tag{1}$$

and distance constraints enforced by the rigid bars c_4, c_5, c_6 hold for 3 other pairs

$$\|M_k\ \mathbf{p}_{k,\ell} + T_k - M_{k'}\ \mathbf{p}_{k',\ell'} + T_{k'}\| = d_{k,\ell,k',\ell'}. \tag{2}$$

2.1 The Unoptimized Polynomial System of Problem `tetra`

An unoptimized *polynomial* system of Problem `tetra` can then be obtained in the following three steps.

1 Observe that all points are covered by $c := \{c_1, c_2, c_3\}$.
2 Pick the home coordinate system $h := c_1$.
3 Resolve c_2 and c_3 in the coordinates of h by solving the 3×3 incidence equations (note that c_1 need not be transformed)

$$\begin{aligned} \mathbf{p}_{1,\beta} &= M_2\mathbf{p}_{2,\alpha} + T_2, \\ M_2\mathbf{p}_{2,\beta} + T_2 &= M_3\mathbf{p}_{3,\alpha} + T_3, \\ M_3\mathbf{p}_{3,\beta} + T_3 &= \mathbf{p}_{1,\alpha} \end{aligned} \tag{3}$$

and three scalar distance equations

$$\begin{aligned} \|\mathbf{p}_{1,\gamma} - (M_2\mathbf{p}_{2,\delta} + T_2)\|^2 &= d_1^2, \\ \|M_2\mathbf{p}_{2,\gamma} + T_2 - (M_3\mathbf{p}_{3,\delta} + T_3)\|^2 &= d_2^2, \\ \|M_3\mathbf{p}_{3,\gamma} + T_3 - \mathbf{p}_{1,\delta}\|^2 \qquad &= d_3^2, \end{aligned}$$

where again the $\mathbf{p}_{i,j}$'s and the distances d_i are given by the fixed lengths within the rigid tetrahedra and the bars. The triple product of rotation matrices and the translation vector we seek to determine are respectively

$$M_i := \begin{bmatrix} \mathsf{s}_{i2}\mathsf{s}_{i3} & \mathsf{s}_{i2}\mathsf{c}_{i3} & -\mathsf{c}_{i2} \\ \mathsf{s}_{i1}\mathsf{c}_{i2}\mathsf{s}_{i3}-\mathsf{c}_{i1}\mathsf{c}_{i3} & \mathsf{c}_{i1}\mathsf{s}_{i3}+\mathsf{s}_{i1}\mathsf{c}_{i2}\mathsf{c}_{i3} & \mathsf{s}_{i1}\mathsf{s}_{i2} \\ \mathsf{s}_{i1}\mathsf{c}_{i3}+\mathsf{c}_{i1}\mathsf{c}_{i2}\mathsf{s}_{i3} & -\mathsf{s}_{i1}\mathsf{s}_{i3}+\mathsf{c}_{i1}\mathsf{c}_{i2}\mathsf{c}_{i3} & \mathsf{c}_{i1}\mathsf{s}_{i2} \end{bmatrix}, \quad T_{c_2} := \left[\begin{smallmatrix} p_i \\ q_i \\ r_i \end{smallmatrix}\right].$$

Here $\mathsf{s}_{i,j}, \mathsf{c}_{i,j}$, $j = 1, 2, 3$ represent the sines and cosines of the rotation angles and are therefore related by the three scalar equations

$$\mathsf{s}_{i,j}^2 + \mathsf{c}_{i,j}^2 = 1.$$

Altogether, we need to solve a system of $n = 18$ polynomial equations in the n variables

$$\mathsf{s}_{ij}, \mathsf{c}_{ij}, p_j, q_j, r_j, \text{ for } 1 \leq i \leq 3, 2 \leq j \leq 3.$$

in order to resolve c. The maximum degree is $d = 6$ since the incidence equations are of degree 3, the distance equations of degree 6 and the trig-relations are of degree 2 in the variables.

2.2 Efficient Pairwise Parametrizations

The first and the third vector equations in (3) can be easily solved for T_2 and T_3 if we pick $\mathbf{p}_{2,\alpha}$ to be the origin in the local coordinates of c_2 and $\mathbf{p}_{3,\alpha}$ to be the local origin of c_3. This leaves a reduced system of 12 equations and unknowns. To further reduce the algebraic complexity, we parametrize the transformations that attach each rigid body to one that has already been assembled in the home coordinate system. Thereby, we eliminate constraint equations and the remaining system in terms of the parameters will be smaller. We use the following notations for points a, b and c (with coordinates $\mathbf{p}_{i,\alpha}$ etc. in the local coordinate system of a rigid body c_i) and abbreviate 'degrees of freedom' as *dofs*.

T_a	translation that maps a to the origin. (T_a^{-1} maps the origin to a.)
R_{ab}	rotation that maps $b - a$ to the x-axis.
M_{bc}	the matrix $[b, c, b \times c] \in \mathbb{R}^{3\times 3}$ whose columns span $\mathbb{R}^3$.
T	undetermined translation (3 dofs).
R	undetermined rotation about the x-axis (1 dof).
Q	undetermined unit quaternion (3 dofs).

In the previous section, we used a parameterization in terms of c, s such that $\mathsf{s}^2 + \mathsf{c}^2 = 1$. And we could have used q_i so that $\sum_{i=0}^{3} q_i^2 = 1$ so that rotations

Table 1: Parameterizations with parameters t_j and their complexity

k	incidences	parameterization in t_j	$w(k)$	degree
3	$(a,b,c) \to (a',b',c')$	$T_{a'}^{-1} M_{b'-a',c'-a'} M_{b-a,c-a}^{-1} T_a$	0	0
2	$(a,b) \to (a',b')$	$T_{a'}^{-1} R_{b'-a'}^{-1} R R_{b-a} T_a$	1	2
1	$(a) \to (a')$	$T_{a'}^{-1} Q T_a$	3	4
0	none	QT	6	4 (1 for T)

Table 1
In the first column, k indicates the number of points shared by two rigid bodies c_i and c'_i. In the second column, any entry q of $(\ldots, q, \ldots) \to (\ldots, q', \ldots)$ indicates that $M_i \mathbf{p}_{i,q} + T_i = M_{i'} \mathbf{p}_{i',q'} + T_{i'}$ needs to be enforced. The third column displays a parameterization in terms of parameters t_j of the transformations T, R and Q. The scalar $w(k)$ in the fourth column counts the number of free parameters (dofs) in the parameterization. The last column lists the maximal degree in the parameters t_j of the rational parametrization.

R and Q are of the form

$$R := \begin{bmatrix} 1 & 0 & 0 \\ 0 & \mathsf{c} & -\mathsf{s} \\ 0 & \mathsf{s} & \mathsf{c} \end{bmatrix}, \quad Q := \begin{bmatrix} 1-2q_2^2-2q_3^2 & 2(q_1q_2+q_0q_3) & 2(q_1q_3-q_0q_2) \\ 2(q_1q_2-q_0q_3) & 1-2q_1^2-2q_3^2 & 2(q_2q_3+q_0q_1) \\ 2(q_1q_3+q_0q_2) & 2(q_2q_3-q_0q_1) & 1-2q_1^2-2q_2^2 \end{bmatrix}.$$

This, however, yields extra equations $\mathsf{s}^2 + \mathsf{c}^2 = 1$ and $\sum_{i=0}^3 q_i^2 = 1$ and hence more than the minimal number of variables. A simple, effective remedy is to parametrize the variables c, s, q_j by stereographic projection:

$$\mathsf{c} := \frac{1-t_0^2}{1+t_0^2}, \quad \mathsf{s} := \frac{2t_0}{1+t_0^2}, \quad q_0 := \frac{1-\sum_{i=1}^3 t_i^2}{1+\sum_{i=1}^3 t_i^2}, \quad q_j := \frac{2t_j}{1+\sum_{i=1}^3 t_i^2}, j = 1,2,3.$$

For Problem `tetra`, if in step **3** we apply the quaternion transformations to c_2 and c_3, we still retain 3×3 incidence equations and three distance equations, as expected reducing the complexity to $n = 12$ polynomial equations in the n variables

$$t_{ij}, p_j, q_j, r_j, \text{ for } 1 \le i \le 3, 2 \le j \le 3.$$

Off hand, the maximum degree is $d = 8$ since the incidence equations are of rational degree 2 over 4 and the distance equations double this count as we clear the denominator to obtain polynomial equations.

2.3 Optimized Choice and Partial Ordering of Pairwise Incidences

Remarkably, the number of unresolved constraints and the degree of the equations encoding Problem `tetra` can be further reduced by (a) looking at the *complete* collection of maximal rigid bodies implied by the constraints in Problem `tetra` and the *all* the lengths or distances that are fixed by the rigid bodies, as in Figure 2 (the *standard collection of rigid bodies* defined in Section 3), and (b) a smart choice and partial ordering of the elimination steps resulting in the parameterized system (the *incidence tree* parametrization defined in Section 4). For this, we observe that the incidence constraints $\mathbf{p}_{i,\beta} = \mathbf{p}_{i+1,\alpha}$ for $i = 1, 2, 3$ and the lengths that are fixed by the tetrahedra together also define a rigid body, namely a triangle. We call this triangle c_4 and add it to

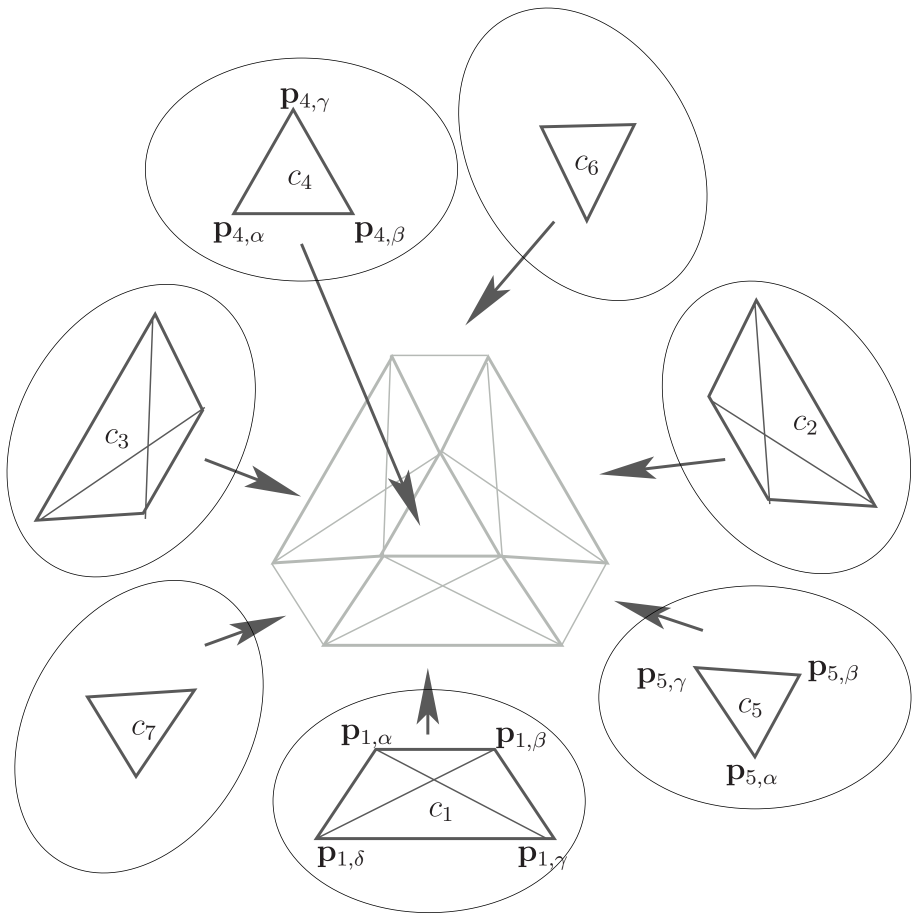

Fig. 2. Problem `tetra`. The new, complete collection of proper-maximal rigid bodies c_i, $i = 1, \ldots, 7$ whose incidences are evident from the central grey graph, such as $\mathbf{p}_{5,\alpha} = \mathbf{p}_{1,\gamma}$ and $\mathbf{p}_{5,\gamma} = \mathbf{p}_{1,\beta}$.

the collection of rigid bodies to be considered (see Figure 2).

1 Choose $c := \{c_1, c_2, c_3, c_4\}$ as the *covering set* of rigid bodies that covers all points. Note that c is not a minimal covering set; i.e., as we know $\{c_1, c_2, c_3\}$ already covers all points and the unoptimized system given earlier was based on this smaller covering set.
2 Pick the home coordinate system $h := c_4$.
3 Since c_4 and each of c_i, $i = 1, 2, 3$ share the two points, $\mathbf{p}_{i,j}$, $j \in \{\alpha, \beta\}$, position and orientation of c_i are now fixed except for rotation about the

axis through the two points. These rotations can be explicitly parameterized so that the incidence equations are resolved and only the three distance equations remain. The complexity of the *parametrized system* is therefore $n = 3$ polynomial equations in the variables t_i for $1 \leq i \leq 3$. The maximum degree is $d = 4$ after clearing the denominator and reduces to 2 due to the identical choice of c_1, c_2 and c_3.

A concrete parametrized system for a specific choice of constraints is shown in Section 5.1, (4).

2.4 Interpretation

The reduction in algebraic complexity from the original to the optimized solution is significant. None of the algebraic and numerical solvers we applied, was able to solve the original problem of size 18×18 or the parameterized problem of size 12×12. By contrast, *all* solutions of the optimized system were obtained in a few seconds using Matlab and Maple. The key question this paper now answers is how to arrive *systematically* at the choice and partial ordering of elimination that yields an optimal parametrized system, in terms of algebraic complexity (number of equations and variables).

In the example, we combinatorially optimize over a natural class of rational parametrized systems that each represent the same well-constrained collection S of incident rigid bodies. In Section 4, we formally such *incidence tree* parametrizations of S. The class uses one of two specific expressions in Table 1 to resolve each pairwise incidence; and each parametrized system in the class corresponds to a subset of rigid bodies that cover all the relevant points that occur in the incidences.

The above example and the example in Figure 3 show that in order for the optimization over the class to be efficient, we have to be careful what collections we resolve. We therefore pose (and answer) the question: *How well does the system of minimal algebraic complexity from the class of incidence tree parametrizations of a collection S compare with parametrizations of S that are not in this class?* This question motivates the formalization – in the next section – of a natural class of *standard collections of rigid bodies S* for which the optimal incidence tree parametrization offers better efficiency and algebraic complexity than any other known method of algebraically solving S. More precisely, we motivate the necessity of each specific property that we use to define a standard collections of rigid bodies by demonstrating that for collections S that do not satisfy that property, there are non-incidence tree parameterizations that have arbitrarily better efficiency and algebraic complexity than the optimal incidence tree parametrization. Fortunately, (a) in most applications naturally occurring collections are standard and (b) good algorithms exist to recursively standardize non-standard collections. Therefore there are efficient ways of solving *any* collection by recursively employing incidence tree parametrizations (Section 4.5).

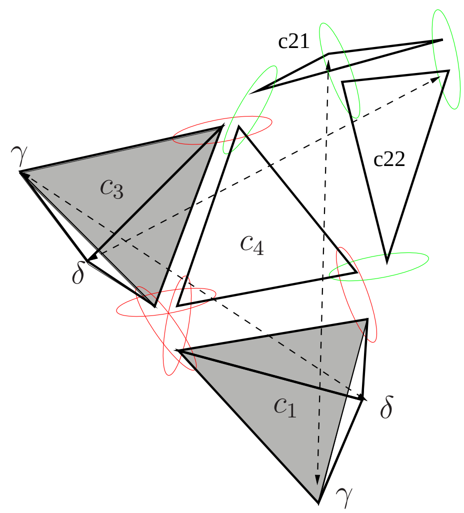

Fig. 3. Problem `tetra` represented in terms of five rigid bodies c_1, c_3, c_4, c_{21}, and c_{22}. These rigid bodies do not form a standard collection of rigid bodies. The intuitive solutions (e.g. attach c_1 and c_3 to c_4, use a quaternion constraint to join c_{21} to c_4 and attach c_{22} to c_{21}) yield a parametrized system of size 6 by 6. The optimal solution (consider c_5 and c_6 as in Figure 2, then attach c_1 and c_3 to c_4 and c_5 to c_1, c_6 to c_3) yields a parametrized system of size 4 by 4. This is worse than the 3 by 3 system obtained by first resolving c_{21} and c_{22} into c_2 with the help of the distance constraint implied by c_4.

3 Standard collections of rigid bodies and the Overlap Graph

To start, we formalize the incidence of rigid bodies.

Definition 1 (Collection of rigid bodies, Covering set) *Let* $c := \{c_1, \ldots, c_n\}$ *be a set of rigid bodies and* X *a set of shared (coordinate free) points in* $\mathbb{R}^3$ *that imply incidences, i.e. points that occur in two or more of the rigid bodies. The pair* (X, c) *is a* valid collection of rigid bodies *(short:* collection*), if for* $i \neq j$*,* c_i *contains at least one point in* X *not in* c_j*, and* $c_i \subsetneq X$*. We call* $c' \subset c$ *a* covering set *of* X *if every point in* X *lies in at least one rigid body in* c'*.*

Essentially, we identify each rigid body c_i with a distinct rigid, proper subset of X. We may assume that

- all distances between any pair of points in a rigid body are fixed and
- the bodies are constrained with respect to each other only by incidence constraints.

The following properties hold for typical collections and are automatic when the collections arise from a good decomposition-recombination process, see Section 4.5. The first, *minimality of the collection*, states that a subset of two or more of the rigid bodies in the collection do not together form rigid body, unless their union includes all the points in X. This turns out to be equivalent to the requirement that no rigid body in the collection can be extended without including all points in X. We then call the rigid bodies in the collection *proper-maximal*. Example 4 in Section 4 shows that inclusion of rigid bodies that are not proper-maximal yields suboptimal incidence trees and hence suboptimal parametrized systems.

Definition 2 (Proper-maximal) *A rigid body c_i is* proper-maximal *in X if there is no subset u of X with $c_i \subsetneq u \subsetneq X$ that represents a rigid body. That is, no such u is rigidified by the incidences within u and the fixed distances within the bodies outside u.*

Example 2 (Proper-maximality) *The subcollection $c_1, c_2, c_3, c_5, c_6, c_7$ (leaving out c_4) in Figure 2 consists of six proper-maximal rigid bodies even though the subcollection is rigid. The rigidity seems to contradict proper-maximality of the six rigid bodies. However, the subcollection includes* all points *in X, hence does not violate the proper-maximality of any of its constituent six rigid bodies.*

Completeness is a second natural property (again guaranteed, for example, by a good decomposition-recombination plan). Completeness requires that *all* proper-maximal rigid sets of points in X are listed in the standard collection of rigid bodies. For example, completeness ensures that rigid body c_4 appears in the standard collection of rigid bodies of Problem `tetra` in Section 2.3, so it can be chosen as the home rigid body in the optimal elimination. With the above definitions, we are ready to define a standard collections of rigid bodies as a complete collection of proper-maximal rigid sets.

Definition 3 (Standard collection of rigid bodies) *The collection (X, c) is a* standard collection of rigid bodies *if the following hold.*

(i) *No pair $c_i \neq c_j$ intersects in more than two points.*
(ii) **(c is complete in X)** *All proper-maximal rigid bodies are in c.*
(iii) *All c_i are proper-maximal.*

Example 3 (Standard collection) *The collection in Figure 2 is standard because each rigid body, $c_1, \ldots, c_7$, is proper-maximal and every proper-maximal rigid body is in the collection.*

Condition (i) reflects the fact that two rigid bodies overlapping on 3 points generically form a rigid body in three dimensions. Proper-maximality (iii) excludes any pair of rigid bodies overlapping on 3 points, unless they form a covering pair and their union includes all points in X. We may exclude that remaining case of triple incidence since such a covering pair of rigid bodies can be instantaneously resolved without solving any system. Condition (i) restricts the number of rigid bodies in c be $O(|X|^3)$ rather than an exponential in $|X|$.

The key structure to efficiently resolve standard collections of rigid bodies is the overlap graph (see for example Figure 6, *left*).

Definition 4 (overlap graph) *An overlap graph $\mathcal{G}(X, c)$ of a standard collection of rigid bodies (X, c) is a weighted undirected graph whose vertices are the rigid bodies c_j in c and whose edges represent incidences between pairs of rigid bodies. If an edge between a pair (c_i, c_j) represents k incidences, the weight $w(k)$ assigned to the edge is the number of remaining dofs of c_j after fixing c_i's position and orientation and resolving the k incidences between them.*

Since we need one parameter to describe the position and orientation of a rigid body with respect to another if the two share 2 points, three if they share one point and six if they share no point, the edge weight matches exactly $w(k)$ according to Table 1:

$$w(0) = 6, \quad w(1) = 3, \quad w(2) = 1.$$

We note that an overlap graph is not a full incidence (hyper)graph since overlap is based on only pairwise incidences. Information such as the same point in X being shared by 3 or more rigid bodies in c is not needed.

4 The class of incidence tree parametrizations for standard collections of rigid bodies

Let us consider *all* possible covering sets defined by Definition 3 (ii). Each covering set $S(c)$ of a collection (X, c) induces a subgraph of the overlap graph $\mathcal{G}(X, c)$. This subgraph contains information about all incidences in the collection c (Figure 4) missing only the information about those constraints, listed as $\mathcal{E}(c \backslash S(c))$, that are additionally required to rigidify the rigid bodies in $c \setminus S(c)$. Since (X, c) is generically well-constrained, this is necessary and sufficient to rigidify the overall collection.

Let $\mathcal{E}(S(c))$ be the incidence constraints of the covering set $S(c)$. An *incidence tree* in $S(c)$ is a spanning tree $\mathcal{T} := \mathcal{T}(S(c))$ of the subgraph (of the overlap graph) induced by $S(c)$. It represents a set $\mathcal{E}(\mathcal{T})$ of constraints that includes:

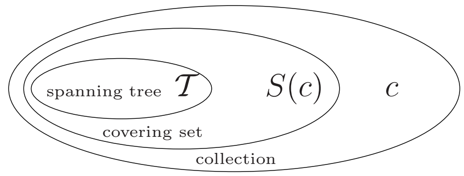

Fig. 4. The goal is to minimize the algebraic complexity of the parametrized system $\mathcal{E}(c\backslash\mathcal{T})$ in terms of the parameters $\{t_j\}$ of Table 1.

(i) a subset of the incidence constraints in $\mathcal{E}(S(c))$ and (ii) the constraints that rigidify the bodies in $S(c)$. We will eliminate the constraints in $\mathcal{E}(\mathcal{T})$ using the parameterizations of Table 1 for $w(k) \in \{1, 3\}$. The remaining constraints in $\mathcal{E}(S(c))$ that have not been eliminated, $\mathcal{E}(S(c)\backslash\mathcal{T})$, together with further distance constraints that are needed to rigidify the bodies that are not in the covering set, $\mathcal{E}(c\backslash S(c))$, form a *parametrized system*, denoted $\mathcal{E}(c\backslash\mathcal{T})$, of some m independent constraints (see Section 4.5 and Figure 4). Since c is assumed to be well-constrained, m equals n, the number of unknowns of the parametrized system. For collections (X, c) in three dimensions, n is at most $3|X| - 6$.

The choice of distances to be put into $\mathcal{E}(c\backslash S(c))$ is not unique. For example, for each of the rigid bodies in $c \setminus S(c)$, the complete graph of pairwise distances is available since each body is rigid. From these, a minimal set is chosen so that $\mathcal{E}(c\backslash S(c)) \cup \mathcal{E}(\mathcal{T})$ rigidifies each of the bodies in $c \setminus S(c)$. Since (X, c) is generically well-constrained, this is necessary and sufficient to rigidify the overall collection. The choice of $S(c)$ and $\mathcal{T}(S(c))$ thus gives a partial order of elimination steps and defines a parametrized system of remaining constraints.

Definition 5 (Parameterizations to be optimized) *The set of all possible choices $S(c)$ and $\mathcal{T}(S(c))$ is the* class of incidence tree parametrizations *for standard collections of rigid bodies.*

Optimizing over this class is equivalent to selecting an appropriate covering set and a spanning tree for the covering set that result in a parametrized system that minimizes algebraic complexity. To select the function to optimize, we observe that all commonly used polynomial system solvers suited to sparse geometric constraint systems take time exponential in the number of variables n (see e.g. [2,20,1,19]) and typically n dominates the algebraic complexity compared to the degree of the equations. By our choice of parameterizations of Table 1, $n = \sum_{e \in \mathcal{T}} w_e(k)$ is the total edge weight of the tree $\mathcal{T}$. An *optimal incidence tree* is therefore a tree of minimum total weight *over all* covering sets.

We can now illustrate the importance of the requirement (iii) of Definition 3.

Example 4 (Non-maximality) *Figure 5 shows two rigid bodies, c_{21} and c_{22}, that are not proper-maximal. The resulting overlap graph has a minimal incidence tree yielding 4 equations in 4 variables (as may be checked by ap-*

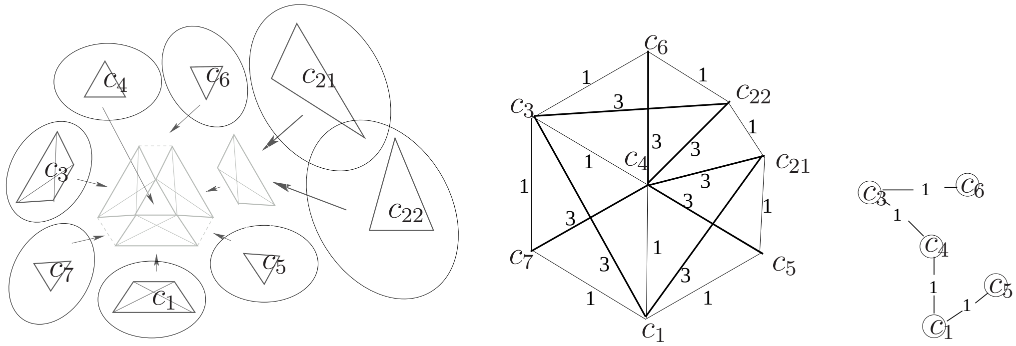

Fig. 5. Problem `tetra`. (*left*) The proper-maximal rigid body c_2 is split into two rigid bodies c_{21} and c_{22}. The resulting 8 rigid bodies no longer represent a standard collection of rigid bodies. (*middle*) The overlap graph of the non-maximal hence non-standard collection of rigid bodies. (*right*) The best incidence tree for this collection has weight 4 rather than the minimal weight 3 for the corresponding standard collection of rigid bodies, as derived in Section 2.3, **3**.

plying the Optimized Incidence Tree Parametrization Algorithm in the next Section 4.1). If the collection is first standardized, (as it would in a canonical decomposition-recombination process, Section 4.5), the optimal incidence tree parametrization can be applied to the standard collection of rigid bodies at each level.

At the first level, the standard collection of rigid bodies consists of three rigid bodies: c_{21}, c_{22} *and the distance constraint between them that is implied by* c_4*. The optimal incidence tree is based on a covering set of* c_{21} *and* c_{22} *and has weight* $w = 1$*. That is, the incidence between* c_{21} *and* c_{22} *is resolved using exactly one variable, for the rotation of* c_{22} *about the line (connecting the 2 points) shared by* c_{21} *and* c_{22} *and the parametrized system consists of one distance constraint between these two bodies.*

At the second level, the standard collection of rigid bodies then has the 7 rigid bodies exactly as shown in Figure 2. As Section 2.3 showed, this level can be resolved by solving 3 equations in 3 variables. Indeed, the optimal incidence tree for this standard collection of rigid bodies has weight 3.

4.1 Optimized Incidence Tree Parametrization Algorithm

The algorithm consists of two parts.

Part 1: Find the optimal incidence tree
Input: A standard collection of rigid bodies (X, c).
Output: An optimal incidence tree $\mathcal{T}(X, c)$.

(1) Over all covering sets $S(c)$, determine the set $\{\mathcal{T}_\ell\}$ of spanning trees (of

the subgraphs of the overlap graph $\mathcal{G}(X, c)$ induced by $S(c)$) of minimum weight: these are candidates for the optimal incidence tree.

(2) Over all choices of trees in $\{\mathcal{T}_\ell\}$ *and roots* determine a rooted tree that minimizes the sum of the depths of all nodes.

Examples of standard collections of rigid bodies and their incidence trees are shown in Figures 6, 11 and 15.

Given
— the optimal incidence tree $\mathcal{T} := \mathcal{T}(X, c)$ and its constraints $\mathcal{E}(\mathcal{T})$ and
— the list $\mathcal{E}(c \backslash \mathcal{T}) := \mathcal{E}(c \backslash S(c)) \cup \mathcal{E}(S(c) \backslash \mathcal{T})$ of incidences and
implied distances that are within the rigid bodies of the original system, but not in the tree $\mathcal{T}$,
the parametrized system is now generated as follows.

Part 2: Elimination and the parametrized system
Input: Constraints $\mathcal{E}(\mathcal{T})$ and $\mathcal{E}(c \backslash \mathcal{T})$ and a tree $\mathcal{T}$ with weights $w(k)$.
Output: The parametrized system $\mathcal{E}_t$ equivalent to $\mathcal{E}(c \backslash \mathcal{T})$, but in terms of the parameters t_j of Table 1.

(1) Traverse $\mathcal{T}$ in *reverse* breadth-first order.
(2) At each node, express all child nodes (and all their children which are already expressed in the child's coordinate system) in the node's coordinate system using the parameterizations of Table 1 corresponding to the weight $w(k)$ of the edge between the node and its child.
(3) At the root, after visiting all nodes, replace the coordinates $\mathbf{p}_{i,j}$ in $\mathcal{E}(c \backslash \mathcal{T})$ by the parameterized coordinates to obtain $\mathcal{E}_t$.

In the final step, the parametrized system is solved by a general algebraic-numeric solver such as [3,6,12].

A subtle point remains discussion. Note that $\mathcal{E}(\mathcal{T})$ and $\mathcal{E}(S(c) \setminus \mathcal{T})$ are essentially fixed by Part 1 (the choice of $S(c)$ and $\mathcal{T}$) of the Optimized Incidence Tree Parametrization Algorithm. However, as input to Part 2 of the algorithm, as pointed out earlier, there are several possible choices of implied distance constraints in $\mathcal{E}(c \setminus S(c))$. It is immediate that this choice cannot affect the number of variables and equations, nor the degree of each variable in the parametrized system. Also the total degree of the system is unaffected by the choice since this choice has to minimally rigidify each of the bodies in $\mathcal{E}(c \setminus S(c))$.

The decrease in the *number of variables and equations* obtained through the optimization is easy to quantify.

Variables and Equations. Without optimization, a well-constrained set c of k rigid bodies forming a covering set results in a system with $\tilde{n} = 6(k-1)$ variables, independent of the number of nonempty overlaps between rigid bodies in the covering set. Using optimization on the other hand, we obtain

$$n = \sum_{j=0}^{3} w(j)k_j = 3k_1 + k_2 + 6(k-1-k_1-k_2) = \tilde{n} - 3k_1 - 5k_2$$

where k_j is the number of pairs of rigid bodies with exactly j incidences in the selected minimum weight spanning tree. Since the tree edges all represent nonempty overlaps of at least one point, optimization at least halves the number of variables and equations (if $k_1 = k-1$) and, at best, reduces it by a factor of six (if $k_2 = k-1$).

Degree. The parametrized system can consist of point-matching constraints and distance constraints. In the unoptimized system, the *total degree* of all distance equations is $d = 6$. In the optimized system, the *total degree* increases with the depth of the spanning tree. However, the *coordinate degree* of the numerator and of the denominator of the constraint systems is at most 4 throughout the elimination since we introduce new variables with each step of the partial elimination. (The coordinate degree is an n-tuple listing the degrees for each variable, e.g. (2,2) for x^2y^2.)

4.3 Overall complexity of resolving the collection

We recall that the overall process of resolving the input collection of rigid bodies consists of two phases. This paper gives a purely combinatorial algorithm for the first phase, i.e., optimizing the algebraic complexity of the system that is input to the second phase, the algebraic-numeric solver. The overall complexity of resolving the collection of rigid bodies is overwhelmingly dominated by the second, algebraic phase that is at least exponential in the size of the parametrized system fed to it. In practice, even systems of size 10 by 10 are very difficult to solve when all solutions are required, as in our problem. We saw above that even in the worst case, our combinatorial optimization halves the size of the parametrized system input to the second phase (in the best case it improves by a factor of 6) while in the unoptimized case, the size of the system is proportional to the size of the input collection of rigid bodies. Hence the first step improves the complexity of the overall process by a factor that is exponential in the size of the input collection. So even for small collections of size 10-30, the first phase yields a huge gain in overall complexity. Asymptotic complexity analysis of the first phase, i.e., obtaining the optimal incidence tree parametrization, is relatively meaningless, for such small collections since even a brute force method is fast and effective. For completeness, we now give the asymptotic complexity of the first phase.

4.4 Complexity of Optimized Incidence Tree Parametrization Algorithm

Recall that obtaining the optimal incidence tree requires finding spanning trees of minimum total weight among all subgraphs induced by covering sets S of the overlap graph $\mathcal{G}(X, c)$. This is a *set cover* problem, with modified objective function, for the special case where the sets are forced to have bounded intersection (of size 1 or 2). The *set cover* problem (where the size of the cover has to be minimized) is NP-complete even when the sets are forced to have bounded intersection. Fortunately, the *set cover* problem has a good polynomial time approximation algorithm based on the primal-dual method in linear programming. Furthermore, our overlap graphs are special since they correspond to standard collections of rigid bodies. Due to maximality and completeness of standard collections of rigid bodies, many types of subgraphs and minors are forbidden in the overlap graph of standard collections of rigid bodies. As a result, the set of candidate incidence trees can be pruned and optimal incidence trees can be found in time polynomial in $|c|$ using the data structures of [5] that efficiently represent and maintain the entire set of spanning forests. (As pointed out earlier, $|c| \leq O(|X|^3)$ for standard collections of rigid bodies, (X, c)). The details of this data structure and analysis are outside the scope and emphasis of this paper.

4.5 Decomposition-Recombination, Hierarchies of standard collections of rigid bodies and Stability

In full generality, our optimization problem is best approached by a recursive or hierarchical decomposition of the original generically rigid system of incidences so that the *subsystems* correspond to rigid subcollections that can themselves be recursively decomposed. In reverse, rigid subsystems can be recombined into a parent rigid body by solving a recombination system. Selecting one instance of each resolved child at a time, we can recombine ever larger parent rigid bodies until the global system is solved. Since the cost of solving any of these recombination systems is dominated by its number of variables n, it is best to recursively decompose in such a manner that the algebraic complexity of the recombination at any given stage is minimized. This is called the optimal *decomposition - recombination (DR)* planning problem [8,9], and is, in general, NP-hard [11]. However, for generically rigid systems, a canonical graph-theoretical decomposition with many nice properties has been developed [15,18], so that, in this paper, we could focus on the solution process at a *single recombination level*. Crucially, since each level is part of a larger, canonical DR process, we need not consider all possible collections of rigid bodies, but only DR collections as in [15,11,18], and these are *standard collections of rigid bodies* (Section 3, Definition 3).

Two very special scenarios merit discussion. (1) *Implied rigid bodies.* In any level of a typical DR process, the standard collection of rigid bodies could contain c_i whose rigidity is implied by other rigid bodies in the collection (in addition to the *inherently* rigid bodies that are rigidified by internal constraints). If such an implied c_i is in the optimal incidence tree then the rigid bodies that imply the rigidity of c_i also have to be resolved even if they are not in the optimal incidence tree. (2) *Stability of optimal incidence tree parameterizations.* Consider the leftmost incidence tree of Figure 11. The set c is well-constrained. There are two non-tree (dashed) edges representing incidences. Choosing the incidence between c_2 and c_3 yields a solvable system. However, choosing the incidence between c_1 and c_4 results in three *dependent* constraints. This stability or independence problem for incidences is addressed and solved in [16] so that we may assume that any choice of the appropriate number of constraints yields a valid choice.

5 Computed Examples

We now solve several geometric instances of Problem `tetra` (Figures 1, 2, 6), our running example; of Problem `pent` (Figure 10) used to illustrate stability issues discussed at the end of Section 4.5; and of Problem `quad` (Figure 14) which serves to illustrate a complete one-parameter family of solutions. We solve the instances using a pipeline of routines. The first part of the combinatorial optimization phase, i.e., the optimal incidence tree algorithm, was implemented as a C++ program. The second part, i.e., the elimination and parameterization, was implemented as a Maple program. The algebraic phase, the solution of the parametrized system, was implemented in Matlab and uses a subdivision-based numerical solver [6]. None of Problem `tetra`, Problem `pent` and Problem `quad` could be solved in their unoptimized form. All three problems are solved in their optimized form with none of the routines in the pipeline taking more than a few seconds.

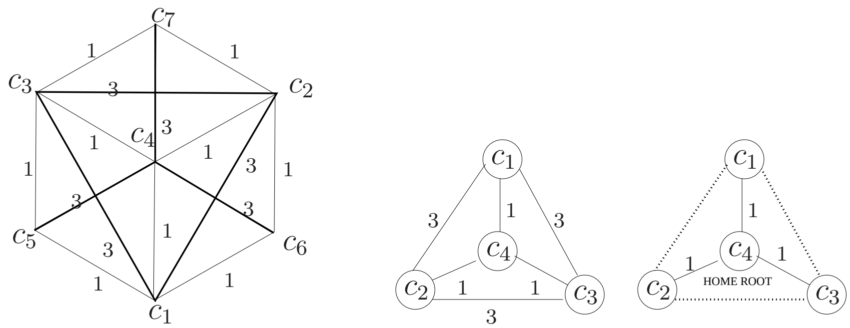

Fig. 6. Problem `tetra` of Figures 1 and 2. (*left*) The overlap graph. (*middle*) The subgraph of the weighted overlap graph induced by one covering set $S(c)$ (the points of c_5 are covered by c_2, c_4 and c_6; the edges of weight 6 representing empty overlaps are omitted). (*right*) The optimal incidence tree $\mathcal{T}$ returned by the Optimized Incidence Tree Parametrization Algorithm has weight 3. There are no non-tree incidences in $\mathcal{E}(S(c) \setminus \mathcal{T})$. The *dotted* lines are the 3 implied distance constraints that form the parametrized system. They constitute $\mathcal{E}(c \setminus S(c))$, and together with the implied distance constraints in $\mathcal{E}(\mathcal{T})$, they rigidify the remaining rigid bodies c_5, c_6, c_7 in $c \setminus S(c)$.

5.1 *Problem* `tetra`

Using the Optimized Incidence Tree Parametrization Algorithm, Problem `tetra` reduces to $n = 3$ equations (see (4) below and Section 2.3). To simplify the visualization of the output, we consider symmetric data. The three rigid bodies c_1, c_2 and c_3 are initialized with *identical local coordinates* of a regular tetrahedron

$$\mathbf{p}_{i,\alpha} := \left[\begin{smallmatrix}-1\\-1\\-1\end{smallmatrix}\right], \quad \mathbf{p}_{i,\beta} := \left[\begin{smallmatrix}-1\\1\\1\end{smallmatrix}\right], \quad \mathbf{p}_{i,\gamma} := \left[\begin{smallmatrix}1\\-1\\1\end{smallmatrix}\right], \quad \mathbf{p}_{i,\delta} := \left[\begin{smallmatrix}1\\1\\-1\end{smallmatrix}\right], \qquad i = 1, 2, 3.$$

The incidence and distance constraints that force the tetrahedra into a specific spatial arrangement are

$$\mathbf{p}_{i,\beta} = \mathbf{p}_{j,\alpha}, \qquad \text{and} \qquad \|\mathbf{p}_{i,\gamma} - \mathbf{p}_{j,\delta}\|^2 = d_i^2,$$

where $j = 1$ if $i = 3$ and $j = i + 1$ otherwise. The underlying geometry is visualized in Figure 7 for one, extreme choice of the distances d_i. We consider three instances.

Problem **`tetra i`.** The maximal value for $d_1 = d_2 = d_3$ that still allows for a real solution is $\bar{d}^2 := 22 + 4\sqrt{2}\sqrt{3}$. Then there is exactly one solution as shown in Figure 7. We juxtapose the spatial view with a more abstract diagram displaying the home rigid body and the parametrized distance constraints since rigid bodies typically obscure one another. After scaling by $(1-t_i)^2 t_i^2 (1-t_j)^2 t_j^2$,

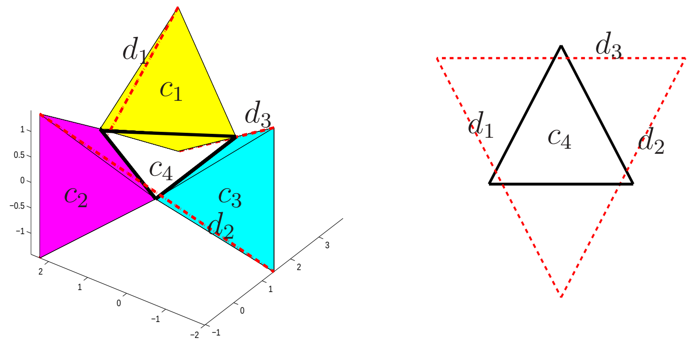

Fig. 7. Problem `tetra i` has a unique solution when the distances (represented by *dashed* line segments) are chosen equal and maximal. (*left*) Geometry of the tetrahedra configuration. The home rigid body c_4 (emphasized triangle in the center) consists of the three line-segments about which the tetrahedral rigid bodies c_1, c_2, c_3 can rotate if no other constraints are present than their 3 pairwise incidences. (*right*) Visualization emphasizing the central triangle and displaying the geometric position of the dashed edges representing the 3 parametrized distance constraints. The end points of the dashed lines meet only in the displayed projection, not in $\mathbb{R}^3$.

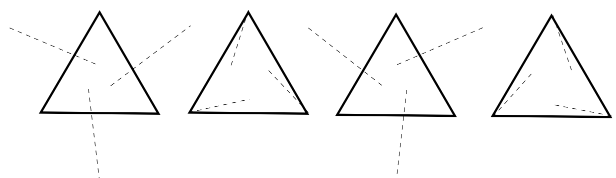

Fig. 8. Problem `tetra ii`. Diagram of the four solutions.

we obtain the following parametrized system for $(i,j) \in \{(1,2),(2,3),(3,1)\}$

$$\begin{aligned}
\Big((16\sqrt{3}+8-16\sqrt{6})t_i^2 + (8\sqrt{6}-8-16\sqrt{3}+8\sqrt{2})t_i \\
+ 4\sqrt{3}-4\sqrt{6}-16-4\sqrt{2}\Big)t_j^2 \\
+\Big((-8\sqrt{2}-8+24\sqrt{6}-16\sqrt{3})t_i^2 + (-24+16\sqrt{3}-16\sqrt{6})t_i \\
+ 4\sqrt{2}+8\sqrt{6}-4\sqrt{3}+32\Big)t_j \\
+(4\sqrt{3}-16-12\sqrt{6}+4\sqrt{2})t_i^2 + (-4\sqrt{3}+32-4\sqrt{2}+8\sqrt{6})t_i \\
= 4\sqrt{6}+16.
\end{aligned} \tag{4}$$

Problem **`tetra ii`.** For $d_1^2 = d_2^2 = d_3^2 = \bar{d}$ (no square), we obtain four solutions with evident symmetries (Figure 8).

Problem **`tetra iii`.** By altering the distances within one rigid body, we obtain ten solutions shown in Figure 9. There are three spatially non-isomorphic configurations.

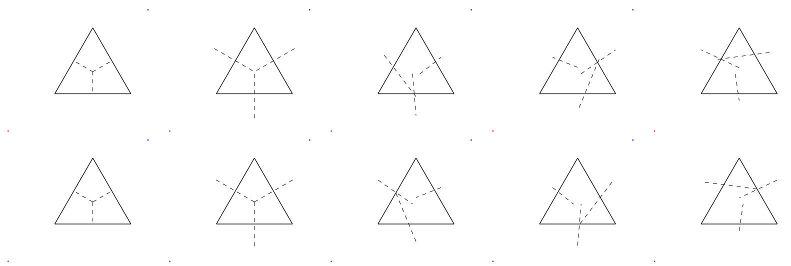

Fig. 9. Problem `tetra` **iii**. Diagram of the ten solutions.

5.2 Problem `pent`

Figure 11 shows covering sets and trees for Problem `pent` in Figure 10 with standard collection of rigid bodies $c := \{c_1, \ldots, c_6\}$. Although the rightmost covering set is smallest, its minimum incidence tree weight of 9 makes it sub-optimal. Any of the middle three covering sets yields trees of weight 5 and the leftmost has the optimal weight 4. Different home (root) rigid bodies, say c_4 versus c_5 for the left-most tree, yield different maximum degrees for the parametrized system.

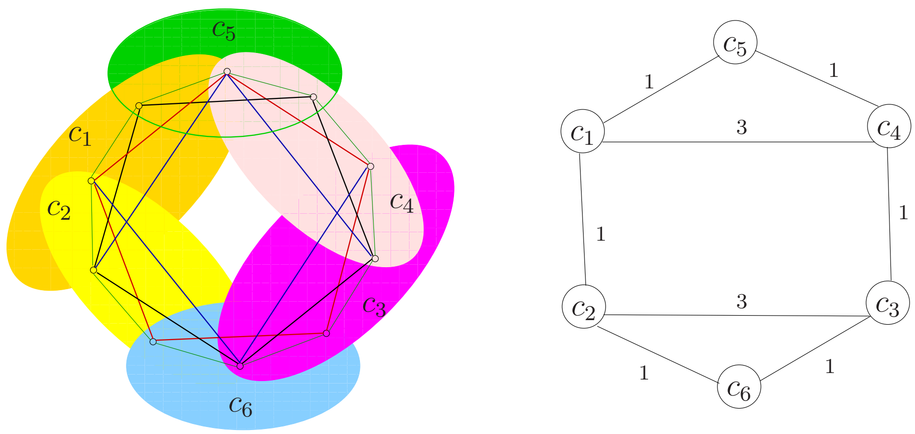

Fig. 10. Problem `pent`. (*left*) Standard collections of rigid bodies and implied internal distances. (*right*) Corresponding weighted overlap graph (edges of weight 6 representing empty overlaps are omitted).

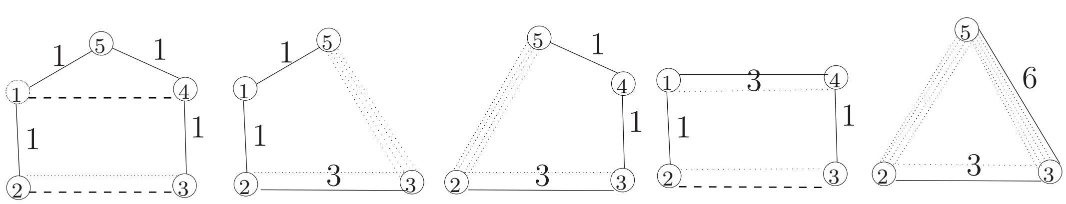

Fig. 11. Problem `pent`. Five spanning trees of covering sets $S(c)$. Implied non-tree distance constraints in $\mathcal{E}(c \setminus S(c))$ are *dotted* - they rigidify the remaining bodies in $c \setminus S(c)$. The non-tree incidences in $\mathcal{E}(S(c) \setminus \mathcal{T})$ (first, fourth tree) are *dashed*. The label i in a circle stands for c_i.

To obtain a concrete instance of Problem `pent`, we define ten (corner) points in the rigid bodies whose x, y coordinates are equally distributed on the unit circle as shown in Figure 10, *left*,

$$\mathbf{q}_i := \left[\begin{smallmatrix}\cos(\theta i)\\ \sin(\theta i)\\ e_i\end{smallmatrix}\right], \text{ where } \theta := 2\pi/10 \text{ and } e_i := \begin{cases}1 & \text{if } i \text{ is odd },\\ 0 & \text{if } i \text{ is even }.\end{cases} \qquad i = 0, 1, \ldots, 9.$$

We initialize the rigid bodies as

$$\begin{aligned}&c_1: \{\mathbf{q}_0, \mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3\}, \quad c_2: \{\mathbf{q}_2, \mathbf{q}_3, \mathbf{q}_4, \mathbf{q}_5\}, \quad c_3: \{\mathbf{q}_5, \mathbf{q}_6, \mathbf{q}_7, \mathbf{q}_8\},\\ &c_4: \{\mathbf{q}_7, \mathbf{q}_8, \mathbf{q}_9, \mathbf{q}_0\}, \quad c_5: \{\mathbf{q}_9, \mathbf{q}_0, \mathbf{q}_1\}.\end{aligned}$$

Optimized Incidence Tree Parametrization Algorithm selects the leftmost minimum spanning tree of Figure 11, with c_5 as the home rigid body. The minimum spanning tree $c_5\{c_1, c_2\}\{c_4, c_3\}$ has two levels: c_2 is a child of c_1 and c_1 is a child of c_5, and c_3 is a child of c_4 and c_4 is a child of c_5. The non-tree edges between c_2 and c_3 represent a stable set [16] of parametrized incidence constraints of degree 2 (the shared point $\mathbf{q}_5$); in addition, the parametrized system includes the distance constraint of degree 4(between $\mathbf{q}_4$ and $\mathbf{q}_6$) as shown in Figure 11, *left*, and Figure 12, *left*. The corresponding four equations between the rigid bodies c_2 and c_3 (three for the shared point $\mathbf{q}_5$ and one distance between $\mathbf{q}_4$ and $\mathbf{q}_6$) force the remaining tetrahedral child rigid bodies c_1, c_2, c_3, c_4 into specific spatial arrangements.

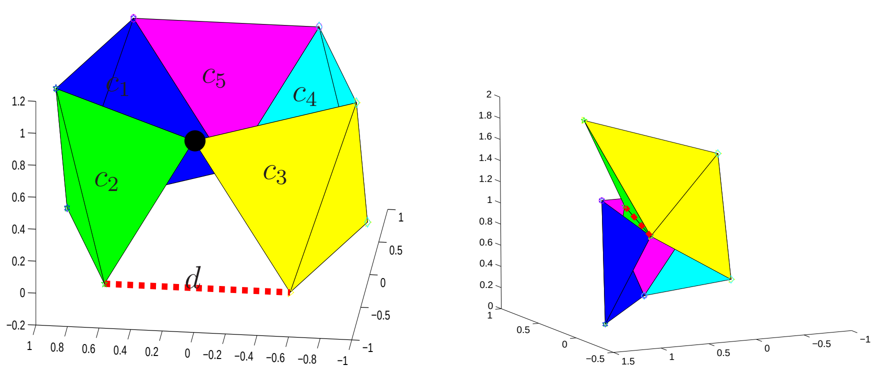

Fig. 12. Problem `pent iv`, two of the eight solutions. For each tetrahedral rigid body, c_1, c_2, c_3, c_4, only two triangle facets are shown. The home rigid body c_5 is a triangle. The dashed line represents the parametrized distance constraint. The parametrized incidence constraint is marked as a black sphere. It forms a triangle with the distance constraint, i.e., together, they constitute the triangle rigid body c_6 of Figure 10 that is not part of the current covering set (see Figure 11, *left*). (*left*) The left configuration can be recognized as representing two pentagons in the z=0 and z=1 plane respectively and rotated by $2\pi/10$ with respect to one another. (*right*) The right realization is self-penetrating.

Problem `pent` **iv.** Let $d = 1.381966011...$, be the square of the length of the edge of a regular unit pentagon. Then there are eight realizations (see Figures 12 and 13). The parametrized system of Problem `pent` does not recommend itself for typesetting.

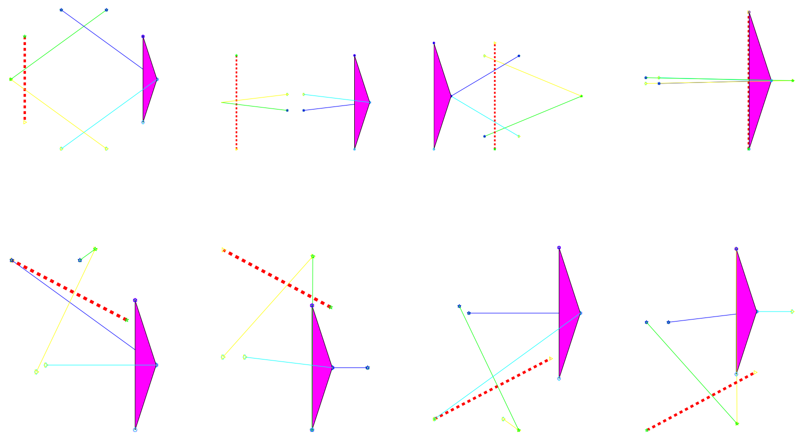

Fig. 13. Diagramatic view (projection) of the eight solutions to the Problem `pent` (**iv**). The *dashed* line represents the parametrized distance constraint. The root rigid body c_5 - a triangle - is displayed, along with four lines connecting the first and last points of each tetrahedral rigid body c_1, c_2, c_3, c_4. The top left diagram corresponds to the left image in Figure 12.

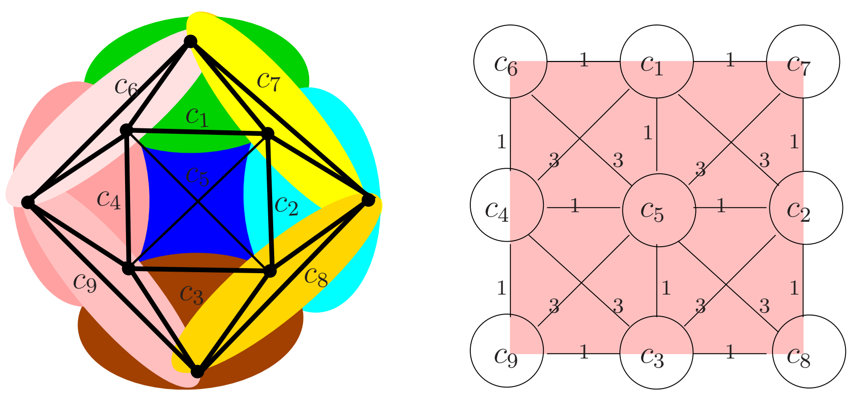

Fig. 14. Problem `quad`. (*left*) Standard collection of 8 rigid bodies (indicated as •) and 18 implied distance constraints within them. The structure is that of a square-base pyramid (see also Figure 16, inset) with the apex split into four (corner) points and the triangular faces folded down. The points are not labeled since they have different labels (and different local coordinates) in different rigid bodies c_i. (*right*) The corresponding weighted overlap graph (all remaining edges of the complete overlap graph are of weight 6 and are omitted).

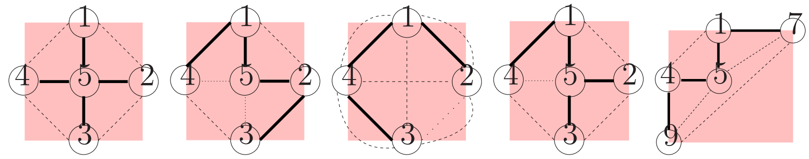

Fig. 15. Problem `quad`. Five spanning trees (*solid* edges) of covering sets (same layout as in Figure 14 *right*; the label i stands for c_i). The sum of weights are from left to right 4,8,9,6 and 4. These numbers match the number of *dashed* distance constraints (= 1 constraint) between rigid bodies and *dotted* incidences that are not tree edges (sharing of nodes = 3 constraints). When optimally rooted (with root c_5), the depth sum of the leftmost tree is 1+1+1+1=4. The depth sum of the rightmost is 2+1+2+1=6. Therefore the leftmost is the optimal incidence tree used for the elimination.

5.3 *Problem* `quad`

The combinatorial structure of Problem `quad` is given in Figure 15. Choosing the minimal covering set (with four rigid bodies, Figure 15, *middle*) does not yield the optimal incidence since the total weight is 9. The minimum weight is 4, and the leftmost tree, with root c_5 and depth 1, is preferred to the rightmost tree with root c_5, sum 4 and depth 2. The Optimized Incidence Tree Parametrization Algorithm automatically generates the parametrized system that fixes the quadrilateral c_5 as home and parametrizes the four points not attached to c_5, each by one parameter. There are four incidence constraints, between c_5 and $c_j, j = 1, 2, 3, 4$, that imply and hence allow to automatically discard, the overlaps between c_i and $c_j, j := i \mod 4 + 1$.

Problem `quad` **v.** The coordinates of the root rigid body c_5 are

$$\mathbf{p}_{5,\alpha} := \left[\begin{smallmatrix}1\\1\\0\end{smallmatrix}\right], \quad \mathbf{p}_{5,\beta} := \left[\begin{smallmatrix}-1\\1\\0\end{smallmatrix}\right], \quad \mathbf{p}_{5,\gamma} := \left[\begin{smallmatrix}-1\\-1\\0\end{smallmatrix}\right], \quad \mathbf{p}_{5,\delta} := \left[\begin{smallmatrix}1\\-1\\0\end{smallmatrix}\right],$$

a 2-unit square. The local coordinates of c_i, $i = 1, 2, 3, 4$ are

$$\mathbf{p}_{i,\alpha} := \left[\begin{smallmatrix}-1\\1\\0\end{smallmatrix}\right], \quad \mathbf{p}_{i,\beta} := \left[\begin{smallmatrix}1\\1\\0\end{smallmatrix}\right], \quad \mathbf{p}_{i,\gamma} := \left[\begin{smallmatrix}0\\4\\0\end{smallmatrix}\right].$$

The four distance constraints are

$$\|\mathbf{p}_{i,\gamma} - \mathbf{p}_{j,\gamma}\| = r. \qquad j := i \mod 4 + 1.$$

For $r = 0$, the four equations of the parametrized system have the particularly compact form

$$0 = -16t_it_j^2 + 16t_j^2t_i^2 + 14t_j^2 - 16t_jt_i^2 + 4t_j - 20t_it_j + 4t_i + 1 + 14t_i^2. \tag{5}$$

Figure 16 displays all possible configurations for different choices of r (with $r = 0$ being the minimum and $r := 2\sqrt{10}$ the maximal distance possible for real solutions).

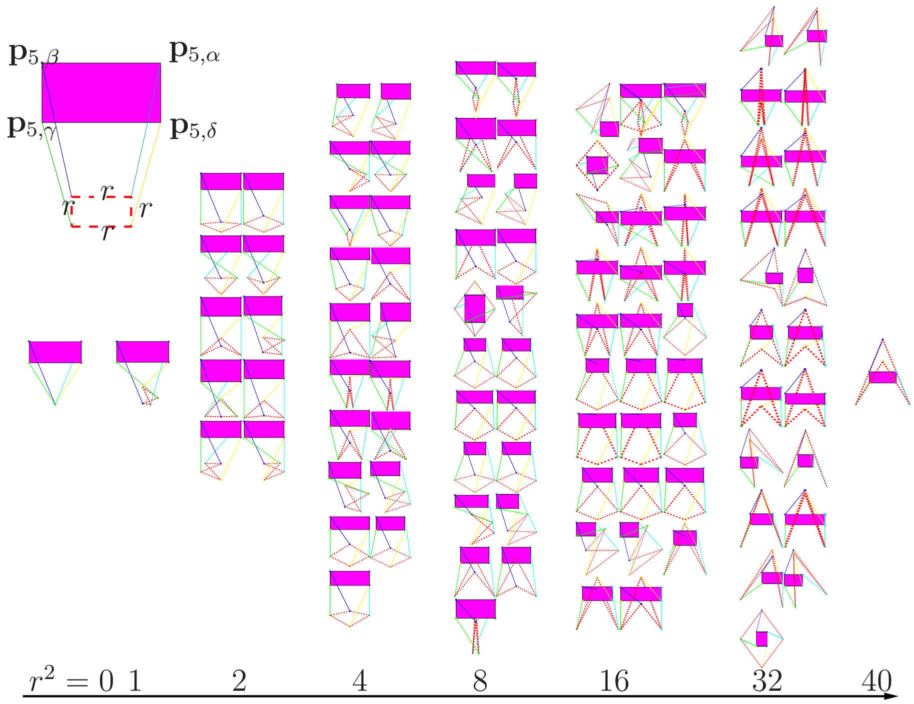

Fig. 16. Problem `quad v`. Possible configurations parametrized by r, the distance of the four distance constraints displayed as *thick, red, dashed* lines. The gallery contains both macro information (number of non-isomorphic solutions) and micro information (geometry of the realization). For example, cf. *upper left* inset, $r = 0$ corresponds to a square-based pyramid (with the solid (*purple*) rectangle representing rigid body c_5; for each of c_1, c_2, c_3 and c_4 only one edge is displayed.) $r^2 = 40$ corresponds to a saddle with two opposing triangle faces flipped up and two flipped down. Find the planar configuration!

6 Discussion

Visualization and Navigation of the Solution Space. The minimum spanning tree $\mathcal{T}$ and the choice of home expose the structure and symmetries of the solution space. Each step on the tree from the leaves to the root corresponds to translating the child's to its parent coordinates and parameterizing the child rigid body by unevaluated rotations. Solving the remaining parametrized system then fixes the rotations in order to enforce the remaining constraints. Thus, the optimized parametrization yields a sequence of key points on a solution path in the combined configuration space of the child rigid bodies. This path has a physical interpretation as a sequence of translations and rotations: Starting from an arbitrary initial position and orientation of the child rigid bodies and ending with their solved position and orientation, i.e., a point in the configuration space of the standard collection of rigid bodies. This path provides an intuitive classification of the exponentially many possible realizations of the collection. By contrast, the large system corresponding to the unoptimized formulation offers few pointers to structure of solution space and is therefore less useful, for example, for solution-space navigation and visual walk-through [17].

Beyond pairwise elimination and incidence trees The focus of this paper is on a relatively small class of incidence tree parametrizations within well-constrained standard collections of rigid bodies. If we expand the class to explicit rational parametrizations that can resolve more general underconstrained subsystems, it may be possible to further reduce the size of the parametrized system. However, optimizing over a significantly larger class is computationally difficult and the expanded class must also be limited in scope since the general problem of finding rational parameterizations of reasonable size leads to the classic hard question of algebraic geometry, whether an ideal/variety has an efficient rational parameterization.

The formalization of standard collections of rigid bodies and optimal incidence tree parametrization hints at parametrizing and reducing the algebraic complexity of more general systems where rigid body incidences are analogous to variables shared by already resolved subsystems in an appropriate decomposition. A natural question is then which algebraic systems can be decomposed into subsystems that resemble standard collections of rigid bodies?

References

[1] Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in Real Algebraic Geometry (Algorithms and Computation in Mathematics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006.

[2] Alicia Dickenstein and Ioannis Emiris Ed.s. *Solving Polynomial Systems, foundations, algorithms and applications, (Algorithms and Computation in Mathematics series number 14)*. Springer-Verlag, 2005.

[3] Gershon Elber. Multivariate (geometric) constraints solving using subdivision based solvers. In Bruno Lévy and Dinesh Manocha, editors, *Proceedings of the 2007 ACM Symposium on Solid and Physical Modeling, Beijing, China, June 4-6, 2007*, page 10. ACM, 2007.

[4] I.Z. Emiris, E. Fritzilas, and D. Manocha. Algebraic algorithms for conformational analysis and docking. *Intern. J. Quantum Chemistry, Special Issue on Symbolic algebra in computational chemistry*, 2006.

[5] David Eppstein. Representing all minimum spanning trees with applications to counting and generation. Technical Report 95-50, Univ. of California, Irvine, Dept. of Information & Computer Science, Irvine, CA, 92697-3425, USA, 1995.

[6] J. Gaukel. *Effiziente Loesung polynomialer und nichtpolynomialer Gleichungssysteme mit Hilfe von Subdivisionsalgorithmen.* PhD thesis, University of Stuttgart, 2003.

[7] Jack E. Graver, Brigitte Servatius, and Herman Servatius. *Combinatorial Rigidity*. Graduate Studies in Math., AMS, 1993.

[8] C. M. Hoffmann, A. Lomonosov, and M. Sitharam. Decomposition of geometric constraints systems, part i: performance measures. *Journal of Symbolic Computation*, 31(4), 2001.

[9] C. M. Hoffmann, A. Lomonosov, and M. Sitharam. Decomposition of geometric constraints systems, part ii: new algorithms. *Journal of Symbolic Computation*, 31(4), 2001.

[10] D.J. Jacobs, A.J. Rader, L.A. Kuhn, and M.F. Thorpe. Protein flexibility predictions using graph theory. *Proteins*, 44(2):150–65, Aug 2001.

[11] Andrew Lomonosov. Graph and Combinatorial Analysis for Geometric Constraint Graphs. Ph.D thesis, Dept. CISE, U of Florida, 2004.

[12] Bernard Mourrain and Jean-Pascal Pavone. Subdivision methods for solving polynomial equations, August 2005.

[13] J. Peters, J. Fan, M. Sitharam, and Y. Zhou. Elimination in generically rigid 3d geometric constraint systems. In *Proceedings of Algebraic Geometry and Geometric Modeling,Nice, 27-29 Sep 2004*, pages 1–16. Springer Verlag, 2005.

[14] M. Raghavan and B. Roth. Inverse kinematics of the general 6R manipulator and related linkages. *J. of Mech. Design, Trans ASME*, 115(3), 1993.

[15] M. Sitharam. Combinatorial approaches to geometric constraint solving: problems, progress, directions. In D. Dutta, R. Janardhan, and M. Smid, editors, *AMS-DIMACS book on computer aided and manufacturing*, volume 67, pages 117–163, 2005.

[16] M. Sitharam. Characterizing well-formed systems of incidences for resolving collecti ons of rigid bodies. *IJCGA*, 16(5-6):591–615, 2006.

[17] M Sitharam, A Arbree, Y Zhou, and N Kohareswaran. Solution management and navigation for 3d geometric constraint systems. *ACM Transactions On Graphics*, 38:22–38, 2006.

[18] M Sitharam and Y Zhou. A tractable, approximate, combinatorial 3d rigidity characterization. *Fifth Automated Deduction in Geometry (ADG)*, 2004.

[19] Sommese and Wampler. *Numerical Solution of Systems of Polynomials arising in Engineering and Science*. World Scientific, 2005.

[20] Bernd Sturmfels. *Solving Polynomial Systems, CBMS series number 97*. American Mathematical Society., 2002.

[21] M.F. Thorpe, M. Lei, A.J. Rader, and L.A. Kuhn. Protein flexibility and dynamics using constraint theory. *J Mol Graph Model*, 19(1):60–69, 2001.