

Characterizing Graphs with Convex and Connected Cayley Configuration Spaces

Meera Sitharam · Heping Gao

Received: 29 September 2008 / Revised: 3 February 2009 / Accepted: 16 March 2009 /

Published online: 3 April 2009

© Springer Science+Business Media, LLC 2009

Abstract We define and study exact, efficient representations of realization spaces *Euclidean Distance Constraint Systems (EDCS)*, which include Linkages and Frameworks. These are graphs with distance assignments on the edges (frameworks) or graphs with distance interval assignments on the edges. Each representation corresponds to a choice of non-edge (squared) distances or Cayley parameters. The set of realizable distance assignments to the chosen parameters yields a parametrized *Cayley configuration space*. Our notion of efficiency is based on the *convexity and connectedness* of the Cayley configuration space, as well as *algebraic complexity of sampling realizations*, i.e., sampling the Cayley configuration space and obtaining a realization from the sample (parametrized) configuration. Significantly, we give *purely graph-theoretic, forbidden minor* characterizations for 2D and 3D EDCS that capture (i) the class of graphs that always admit efficient Cayley configuration spaces and (ii) the possible choices of representation parameters that yield efficient Cayley configuration spaces for a given graph. We show that the easy direction of the 3D characterization extends to arbitrary dimension d and is related to the concept of d -realizability of graphs. Our results automatically yield efficient algorithms for obtaining exact descriptions of the Cayley configuration spaces and for sampling realizations, without missing extreme or boundary realizations. In addition, our results are *tight*: we show counterexamples to obvious extensions.

This is the first step in a systematic and graded program of combinatorial characterizations of efficient Cayley configuration spaces. We discuss several future theoretical and applied research directions.

In particular, the results presented here are the first to completely characterize EDCS that have connected, convex and efficient Cayley configuration spaces, based

Supported in part by NSF Grants EIA 02-18435, CCF 04-04116, and a research gift from SolidWorks.

M. Sitharam (✉) · H. Gao

CISE Department, E301 CSE Building, University of Florida, P.O. Box 116120, Gainesville, FL 32611-6120, USA

e-mail: sitharam@cise.ufl.edu

on precise and formal measures of efficiency. It should be noted that our results *do not rely on genericity* of the EDCS. Some of our proofs employ an unusual interplay of (a) classical analytic results related to positive semi-definiteness of Euclidean distance matrices, with (b) recent forbidden minor characterizations and algorithms related to d -realizability of graphs. We further introduce a novel type of restricted edge contraction or reduction to a graph minor, a “trick” that we anticipate will be useful in other situations.

Keywords Underconstrained Geometric Constraint System · Mechanism · Cayley configuration space · Combinatorial rigidity · Linkage · Framework · Graph Minor · Graph Characterization · Distance geometry · Convex · Semidefinite and Linear programming · Algebraic Complexity

1 Introduction

A *Euclidean Distance Constraint System (EDCS)* (G, δ) is a graph $G = (V, E)$ together with an assignment of distances $\delta(e)$, or distance intervals $[\delta^l(e), \delta^r(e)]$ to the edges $e \in E$. A d -dimensional *realization* is the assignment p of points in \mathbb{R}^d to the vertices in V such that the distance equality (resp., inequality) constraints are satisfied: $\delta(u, v) = \|p(u) - p(v)\|$ (resp., $\delta^l(u, v) \leq \|p(u) - p(v)\| \leq \delta^r(u, v)$). Note the EDCS with distance equality constraints, (G, δ) , is also called a *linkage* and was originally referred to as a *framework* in combinatorial rigidity terminology; more recently a framework (G, p) includes a specific realization p , and the distance assignment δ is read off from p .

We seek *efficient representations of the realization space* of an EDCS. We define a *representation* to be (i) a choice of parameter set, specifically a choice of a set F of non-edges of G , and (ii) a set $\Phi_F^d(G, \delta)$ of possible distance values $\delta^*(f)$ that the non-edges $f \in F \subseteq \overline{E}$ can take while ensuring existence of at least one d -dimensional realization for the augmented EDCS $(G \cup F, \delta(E), \delta^*(F))$. In the presence of inequalities, the Cayley configuration space is denoted $\Phi_F^d(G, [\delta^l, \delta^r])$ and the augmented EDCS is $(G \cup F, [\delta^l(E), \delta^r(E)], \delta^*(F))$. Here $G \cup F$ refers to a graph $H := (V, E \cup F)$. In other words, in this manuscript, our representations are in *Cayley parameters* or non-edge distances: The set $\Phi_F^d(G, \delta)$ (resp., $\Phi_F^d(G, [\delta^l, \delta^r])$) is the *projection* onto the Cayley parameters in F , of the Cayley–Menger semi-algebraic set with fixed (G, δ) (resp., $(G, [\delta^l, \delta^r])$) [6, 8, 29]. This is also the set of d dimensional $|V| \times |V|$ Euclidean distance matrix completions of the partial distance matrix specified by (G, δ) [1].

We refer to the representation $\Phi_F^d(G, \delta)$ (resp., $\Phi_F^d(G, [\delta^l, \delta^r])$) as the *Cayley configuration space* of the EDCS (G, δ) (resp., $(G, [\delta^l, \delta^r])$) *on the parameter set* F of non-edges of G .

Note For ease of exposition, from now on EDCS will generally refer to distance equality constraints only. We will indicate with explicit remarks when a theorem is applicable to EDCS with distance inequality constraints as well.

The Ph.D. thesis [11] formulates the concept of efficient Cayley configuration space description for EDCS by emphasizing the exact choice of parameters used to

represent the realization space. This sets the stage for a *mostly combinatorial, and complexity-graded* program of investigation. An initial sketch of this program was presented in [12]; a comprehensive list of theoretical results and applications to date can be found in the Ph.D. thesis [11].

1.1 Organization

In Sects. 2 and 2.1, we motivate and give a brief background for the overall program of investigation. The questions of interest and contributions of this manuscript are listed in Sect. 3. Their novelty and technical significance are outlined in Sect. 4 together with related work. Formal results and proofs are presented in Sect. 5. We conclude with suggestions for future work in Sect. 6.

2 Motivation

Describing and sampling the realization space of an EDCS is a difficult problem that arises in many classical areas of mathematics and theoretical computer science and has a wide variety of applications in computer aided design for mechanical engineering, robotics and molecular modeling. Especially for *underconstrained or independent and not rigid* EDCS whose realizations have one or more degrees of freedom of motion, progress on this problem has been very limited.

Existing methods for sampling EDCS realization spaces often use Cartesian representations, factoring out the Euclidean group by arbitrarily “pinning” or “grounding” some of the points’ coordinate values. Even when the methods use “internal” representation parameters such as Cayley parameters (non-edges) or angles between unconstrained objects, the choice of these parameters is ad hoc. While Euclidean motions may be automatically factored out in the resulting parametrized Cayley configuration space, for most such parameter choices, the Cayley configuration space is still a topologically complex semi-algebraic set, sometimes of reduced measure in high dimensions.

After the representation parameters are chosen, the method of sampling the Cayley configuration space often reduces to “take a uniform grid sampling and throw away sample configurations that do not satisfy given constraints.” Since even Cayley configuration spaces of full measure (representation using lowest possible number of parameters or dimensions) often have complex boundaries, potentially with cusps and large holes, this type of sampling is likely to miss extreme and boundary configurations and is moreover computationally inefficient. To deal with this, numerical, iterative methods are generally used when the constraints are equalities, and in the case of inequalities, probabilistic “roadmaps” and other general collision avoidance methods are used. They are approximate methods. If the Cayley configuration space is relatively low dimensional, then initial sampling is used to provide an approximate and refinable representation of the Cayley configuration space, using traditional approximation theory methods such as splines or computational geometry representations, for example, based on Voronoi diagrams. Thereafter this approximate representation is used to guide more refined sampling. All of these are approximate methods that do not leverage exact descriptions of the Cayley configuration space.

Two related problems additionally occur in NMR molecular structure determination and wireless sensor network localization: completing a partially specified Euclidean Distance Matrix in a given dimension; and finding a Euclidean Distance Matrix in a given dimension that closely approximates a given Metric Matrix (representing pairwise distances in a metric space) [5, 10]. The latter problem also arises in the study of algorithms for low distortion embedding of metric spaces into Euclidean spaces of fixed dimension [2]. Both of these can be viewed as searching over a Cayley configuration space of an EDCS. But the common methods for these problems are different from those used for exploring Cayley configuration spaces. One reason for this is that usually only one realization is usually sought, which optimizes some appropriately chosen function; the goal is not sampling or description of the entire Cayley configuration space. Common methods for these problems are: (i) either use semi-definite programming, since Euclidean Distance Matrices in a specified dimension are directly related to Gram matrices which are positive semidefinite matrices of a specified rank; (ii) or iteratively enforce the Cayley–Menger determinantal conditions that characterize Euclidean Distance Matrices in a specified dimension.

2.1 Exact, Efficient Cayley Configuration Spaces

Motivated by these applications, our emphasis is on *exact, efficient* description of the Cayley configuration space of *underconstrained* or *independent and not rigid* EDCS. (i) An exact algebraic description guarantees that boundary and extreme configurations are not missed during sampling, which is important for many applications. (ii) An efficient description (i.e., low dimensional, full measure, convex, using few polynomial or even linear inequalities, whose coefficients are obtained efficiently from the given EDCS) is important for tractability of the sampling algorithm.

Efficiency refers to several factors. We list four efficiency factors that are relevant to this manuscript. The first factor is the *sampling complexity*: given the EDCS (G, δ) , (i) the complexity of computing (ia) the set of Cayley parameters or non-edges F and (ib) the description of the Cayley configuration space $\Phi_F^d(G, \delta)$ as a semi-algebraic set, which includes the algebraic complexity of the *coefficients* in the polynomial inequalities that describe the semi-algebraic set, and (ii) the descriptive algebraic complexity, i.e., number, terms, degree, etc. of the polynomial inequalities that describe the semi-algebraic set. These together determine the complexity of sampling or walking through configurations in $\Phi_F^d(G, \delta)$.

Concerning (ia), it is important to note that most choices of Cayley parameters (non-edges) to represent the realization space of (G, δ) give inefficient descriptions of the resulting parametrized Cayley configuration space (see, for example, Figs. 1 and 2). Hence, we place a strong emphasis is on a *systematic, combinatorial choice* of the Cayley parameters that *guarantee* a Cayley configuration space with *all* the efficiency requirements listed here. Further, we are interested in combinatorially characterizing *for which graphs G such a choice even exists*.

The second efficiency factor is the *realization complexity*. Note that the price we pay for insisting on exact and efficient Cayley configuration spaces is that the map from the traditional Cartesian realization space to the parametrized Cayley configuration space is many-one, i.e., *each parametrized or Cayley configuration could correspond to many (but at least one) Cartesian realizations*.

Fig. 1 When parameters for the EDCS in (a) are chosen to be the two dashed non-edges, we get a convex 2D Cayley configuration space shown in (b)





Fig. 2 When parameter for the EDCS in (a) is chosen to be the dashed non-edge, we get a disconnected 2D Cayley configuration space shown in (b). The realization $p(v_1)$ can lie in either of the two solid arc segments of the circle labeled $p(v_1)$, yielding two disconnected intervals for the Cayley configuration space on the non-edge (v_1, v_3) as shown in (c)



However, we circumvent this difficulty by defining and studying realization complexity as one of the *requirements* on efficient Cayley configuration spaces, i.e., we take into account that the realization step typically follows the sampling step, and ensure that one or all of the corresponding Cartesian realizations can be obtained efficiently from a parametrized sample configuration.

A third efficiency factor is *generic completeness*, i.e., we would like (a) each configuration in our parametrized configuration space to generically correspond to at most finitely many Cartesian realizations and (b) we would like the Cayley configuration space to be of *full measure*, and, in particular, they use *at most* as many parameters or dimensions as the internal degrees of freedom of G . Specifically (a) means $G \cup F$ is rigid and (b) means $G \cup F$ is *not overconstrained*, i.e., it is *independent*. This generic completeness means that the graph $G \cup F$ is *well-constrained* i.e., *minimally rigid*.

Note In this paper, unless otherwise specified, we always assume that Cayley configuration spaces have full measure.

A fourth and fifth important efficiency factors are *topological and geometric complexity*, for example, number of connected components, and convexity. Convexity and connectedness are natural properties to study since they facilitate convex programming and other efficient methods for sampling. Another crucial reason for studying convexity is that results (such as those presented here) about convex configuration spaces readily generalize from Euclidean distance equality constraint systems (e.g., frameworks) to Euclidean distance inequality constraint systems (e.g., tensegrities and partially specified Euclidean Distance Matrices with intervals as entries).

2.2 Combinatorial Characterization

Combinatorial characterizations of generic properties of EDCS are the cornerstone of combinatorial rigidity theory. In practice, they are crucial for tractable and efficient geometric constraint solving, since they are used to analyze and decompose the underlying algebraic system. So far such characterizations have been used primarily for broad classifications into well-, over-, and underconstrained, detecting dependent constraints in overconstrained systems and finding completions for underconstrained systems. Such combinatorial characterizations have been missing in the finer classification of underconstrained systems according to the efficiency or complexity of their Cayley configuration space. Our emphasis in this respect is the surprising fact that there is a clean combinatorial characterization *at all* for the algebraic complexity of configuration spaces. This, however, is a necessary step in efficiently decomposing and analyzing underconstrained systems.

3 Questions and Contributions

Next we give 4 natural questions concerning efficient configuration spaces and the contribution of this manuscript towards answering them.

3.1 Question 1: Graphs with Connected, Convex, Linear Polytope 2D Cayley Configuration Spaces

We are interested in characterizing G for which there is a set F of non-edges such that for all distance assignments $\delta(E)$ (resp., intervals $[\delta^l(E), \delta^r(E)]$), the d -dimensional Cayley configuration space $\Phi_F^d(G, \delta)$ (resp., $\Phi_F^d(G, [\delta^l, \delta^r])$), is convex or connected, or is a linear polytope. Furthermore, given G in this characterized class, we would like to characterize the corresponding sets F of non-edges. These are exactly the parameter choices that yield efficient Cayley configuration spaces.

Note Here the phrase *linear polytope* refers to aspect (ib) of the sampling complexity defined in Sect. 2: The coefficients of the linear inequalities that define the bounding hyperplanes of the polytope are simple linear combinations of the given δ determined by G and F . Contrast this with a polytope for which the coefficients of the bounding linear inequalities are the solution of an arbitrary polynomial system determined by (G, δ) and F .

- (1) In Theorem 5.10, we give an *exact* characterization of the class of graphs G all of whose corresponding EDCS (G, δ) admit a 2D (*generically complete*), *linear polytope* Cayley configuration space. The theorem also shows that the characterization remains unchanged if the Cayley configuration space is merely required to be convex, and further if it is merely required to be connected.
- (2) For a graph in the above class, in Theorem 5.11 we give an exact combinatorial characterization of the choices of Cayley parameters (non-edges) that ensure a (*generically complete*), *linear polytope* Cayley configuration space.

- (3) Both above results rely on key Theorem 5.1 (in turn based on Theorem 5.2) that characterizes a graph G along with a non-edge f such that for all distance assignments $\delta(E)$, the 2D Cayley configuration space $\Phi_f^2(G, \delta)$, is a *single interval*. We additionally show in Observation 5.9 that this result is tight in that the obvious analog of this result fails in 3D.
- (4) Observation 5.12 shows that while the forward direction of Theorems 5.10, 5.11, and 5.1 for pure distance constraints holds directly for interval constraints $(G, [\delta^l, \delta^r])$, the reverse direction fails. However, in Theorem 5.13, we give an *exact* characterization of the class of graphs G all of whose corresponding EDCS $(G, [\delta^l, \delta^r])$ admit a 2D (generically complete), linear polytope, convex or connected Cayley configuration space.

We note that the forward direction of the above theorems (that the graph-theoretic property always admits a convex, connected, linear polytope Cayley configuration space) is straightforward. It is the reverse direction that is surprising and the proof requires a new type of restricted edge-contraction reduction to a graph minor (see Sect. 4).

3.2 Question 2: Graphs with Universally Inherent, Connected and Convex Configuration Spaces

One can view Contributions (2) and (3) above as characterizing pairs (G, F) such that G always admits a connected or convex Cayley configuration space on F . Sometimes, it is more convenient to instead characterize the graphs $H = G \cup F$. In particular, we say that a graph H always admits an *inherent* connected or convex Cayley configuration space, if *there exists* a partition of the edges of H into $E \cup F$ so that the graph $G = (V, E)$ always admits a connected or convex Cayley configuration space on F . In other words, for all distance assignments $\delta(E)$ (resp., intervals $[\delta^l(E), \delta^r(E)]$) for the graph $G = (V, E)$, the d -dimensional Cayley configuration space $\Phi_f^d(G, \delta)$ (resp., $\Phi_f^d(G, [\delta^l, \delta^r])$), is connected or convex. We additionally consider the following strong property. We say that a graph H always admits *universally inherent* connected or convex Cayley configuration spaces, if *for every* partition of the edges of H into $E \cup F$, the graph $G = (V, E)$ always admits a connected or convex Cayley configuration space on F . We are interested in combinatorially characterizing graphs H that always admit universally inherent connected or convex Cayley configuration spaces.

- (5) A graph is *d-realizable* if for every δ for which the EDCS (G, δ) has a Euclidean realization in any dimension, it also has a realization in \mathbb{R}^d . This useful notion of d -realizability was introduced by [3, 4], which also showed a forbidden minor characterization of such graphs for $d \leq 3$. For any dimension d , we show in Theorem 5.15 that the class of d -realizable graphs always admit universally inherent connected Cayley configuration spaces, that are in fact convex over *squared* Cayley parameters. We refer to those as *convex squared Cayley configuration spaces*. This result holds also when the corresponding EDCS use distance intervals.
- (6) Theorem 5.16 shows the reverse direction of (5) for 3D, and thus shows that 3-realizable graphs H are *exactly* the ones that always admit universally inherent connected and convex 3D squared Cayley configuration spaces, also when

the corresponding EDCS use distance intervals. Thus, by [3, 4], this class also has forbidden minor characterization. In Observation 5.9, we observe that both Theorems 5.15 and 5.16 are weak statements for 2D—much stronger statements follow directly from (2) above. For example, it follows from (2) that if a graph is not 2-realizable, then there is a natural component of the graph that does not even admit an inherent connected Cayley configuration space description, leave alone a universally inherent one.

3.3 Question 3: Efficiently Recognizing Graphs with Connected and Convex Configuration Spaces

- (7) Both characterizations in Contributions (1) and (6) (Theorems 5.11 and 5.16) directly point to efficient, existing algorithms that *recognize* whether a given graph always admits (universally inherent) connected, convex, (generically complete), linear polytope 2D and 3D (squared) Cayley configuration spaces. For example, a recent algorithm for recognizing 3-realizable graphs [36] was given based on a characterization of such graphs in [3].

3.4 Question 4: Sampling and Realization Complexities

The practical use of the contributions (1)–(6) above becomes apparent when we answer the following question. Given an EDCS (G, δ) (resp., $(G, [\delta^l, \delta^r])$) where G is in one of the classes characterized in contributions (1) to (6), what is the complexity of computing (i) an appropriate set of Cayley parameters or non-edges F (ii) the exact description of $\Phi_F^d(G, \delta)$ (resp., $\Phi_F^d(G, [\delta^l, \delta^r])$) as a semi-algebraic set (iii) a Cartesian realization of a given parametrized configuration in $\Phi_F^d(G, \delta)$. Here (i) and (ii) determine the sampling complexity and (iii) determines realization complexity.

- (8) For 2D, Theorems 5.10 and 5.11 show that the time complexities for (i), (ii) and (iii) are linear. For 3D, our result in Theorem 5.16 only pertains to universally inherent Cayley configuration spaces, and we only have a weak analog of Theorem 5.11; hence, Question 4(i) only marginally applies. We observe that a result from [36] gives a weak answer for question (i) and an $O(|E|^4)$ time complexity for (iii). The complexity for (ii) is an open problem and is discussed in Sect. 6. However, by employing Contribution (5) (Theorem 5.16), the complexity of obtaining *one* configuration is seen to be polynomial (in contrast to obtaining a full description of the Cayley configuration space as a semi-algebraic set).

4 Novelty and Related Work

- The study of the Cartesian realizations of plane linkages or EDCS with distance equality constraints has a long history [25]: Over a century ago, Kempe showed that any plane algebraic curve can be traced by a point in the realizations of a linkage. Furthermore, there are extensive studies of the topology of the configuration space of, for example, polygonal linkages parametrized by Cayley parameters, see, for example, [14].

However, to the best of our knowledge, ours is the first combinatorial or graph-theoretic forbidden minors characterization of geometric and topological complexity such as convexity, and connectedness of Cayley configuration spaces in Cayley parameters even for the plane, and certainly for 3D.

Conceptually, the results on global rigidity [20, 21] (resp., globally linked pairs [22]) are related, since they combinatorially characterize when the Cayley configuration space on all (resp., specific) Cayley parameters (non-edges) is a single point. However, these characterizations hold only generically [9, 17] as is customary for many combinatorial properties related to rigidity. In contrast, we note that our characterizations apply to *all* EDCS' (frameworks) and not just generic frameworks. This is a crucial distinction that is needed to reconcile apparent discrepancies of the two types of results, as we elaborate in Sect. 6.1.2 together with a conjecture concerning the modification of our characterization for generic frameworks.

Moreover, in Sect. 6.1.6 we point out that the class of EDCS with convex (squared) Cayley configuration spaces is much larger than in our characterization, when the distances associated with the edges of the EDCS are known to be special. Furthermore, in Sect. 6.1.4 we point out that the class of EDCS that admit connected 3D Cayley configuration spaces—that are not necessarily universally inherent—is also much larger than in our characterization. Both these observations have computational chemistry applications as pointed out in Sects. 6.2.2 and 6.2.3.

Our results additionally yield a combinatorial characterization of sampling complexity. This incorporates the complexity of (i) obtaining the chosen set of parameters and (ii) the algebraic complexity of obtaining the description of the Cayley configuration space from the given graph. This, in turn, includes the descriptive complexity of the Cayley configuration space as a semi-algebraic set, such as the number and degree of the polynomials. The characterizations moreover incorporate realization complexity, i.e., the complexity of obtaining a realization from a parametrized configuration.

To the best of our knowledge, the only results of a similar flavor are: the result of [31] that shows the equivalence of Tree- or Triangle-decomposability [30] and Quadratic or Radical realizability for planar graphs; and the result characterizing the sampling complexity of 1-dof Henneberg-1 graphs [13].

- Concerning the use of Cayley parameters or non-edges for parametrizing a *generically complete* Cayley configuration space, [34] as well as [24, 37] study how to obtain “completions” of underconstrained graphs G , i.e., a set of non-edges F whose addition makes G minimally rigid or well-constrained. Both are motivated by realization complexity of underconstrained EDCS, i.e., efficiently obtaining a realization *given the parameters values of a configuration*, i.e., once the distance values of the completion edges in F are given. In particular, [24] also guarantees that the completion ensures Tree- or Triangle-decomposability, thereby ensuring low realization complexity. However, they do not even attempt to address the question of how to find realizable distance *values* for the completion edges. Nor do they concern themselves with the geometric, topological or algebraic complexity of *the set of distance values that these completion non-edges can take*, nor the complexity of obtaining a description of this Cayley configuration space, given the EDCS (G, δ) and the non-edges F , nor a combinatorial characterization of graphs

for which this sampling complexity is low. The latter factors, however, are crucial for tractably analyzing and decomposing underconstrained systems and for sampling their Cayley configuration spaces *in order* to obtain the corresponding realizations. The problem has generally been considered too messy, and barring effective heuristics for certain cases, for example in [28], there has been no systematic, formal program to study this problem.

- Some of the proofs (e.g., Theorem 5.15) employ an unusual interplay of (i) classical analytic results related to (squared) Euclidean distance matrices, such as positive semi-definiteness that date back to [33], with (ii) recent graph-theoretic characterizations [3, 4] related to d -realizability. This further permits us to directly apply a recent result about efficient realization of 3-realizable EDCS [36].
- Some of the proofs, e.g., Theorem 5.2, use a novel type of restricted edge-contraction reduction to a graph minor which disallows edge removals a specified pair of vertices to remain distinct and to remain a non-edge. We anticipate that this new “trick” could be useful in other situations [18].

5 Theorems

5.1 Basics

We start with basic definitions and facts. Take two graphs G_1 and G_2 that both contain a complete graph on k vertices, K_k , as a proper subgraph. For any matching of the vertices of the two K_k 's, by identifying the matched pairs, we can get a new graph G_3 . We call this procedure a k -sum of G_1 and G_2 [3]. A graph is a k -tree if it is a k -sum of K_{k+1} 's. Given a graph, we can run the inverse operations of k -sum to get a set of k -sum components. If we cannot run the inverse operations of k -sum for a component, we say that component is a *minimal (k -sum) component*. Given a non-edge f , a *minimal k -sum component containing f* is a minimal subgraph that is both a k -sum component and contains the vertices of f . We emphasize that the phrase does *not* refer to a k -sum component that is minimal and happens to contain f .

A graph is a *partial k -tree* if it is a subgraph of a k -tree. Please refer to Fig. 3 for 2-trees and partial 2-trees and Fig. 4 for 2-sum and 2-sum component. It is not hard to see that partial 2-trees are exactly the 2-realizable graphs. While partial 3-trees are, in fact, 3-realizable, the class of 3-realizable graphs includes non-partial 3-trees as well.

In [3, 4], a useful *forbidden-minor* characterization of such graphs is given. A graph G has a graph K as a *minor* if there is a vertex induced subgraph of G that can be reduced to K via edge removals and edge contractions (coalescing or identifying the 2 vertices of an edge). It is not hard to see that partial 2-trees are exactly the graphs that avoid K_4 minors.

Next we give basic 2D combinatorial rigidity definitions based on Laman's theorem [26]. For 3D, no combinatorial definitions exist. For the corresponding algebraic definitions, please see, for example, [19] (combinatorial rigidity terminology), [15] and [34] (geometric constraint solving terminology).

In 2D, a graph $G = (V, E)$ is *well-constrained* or *minimally rigid* if it satisfies the Laman conditions [26], i.e., $|E| = 2|V| - 3$ and $|E_S| \leq 2|V_S| - 3$ for all subgraphs

Fig. 3 The graph of only *solid edges* is an underconstrained partial 2-tree while the graph of both *solid and dashed edges* is well-constrained and is a 2-tree





Fig. 4 A 2-sum of five minimal 2-sum components (marked by *dashed circles*). The 2-sum component in the middle is a partial 2-tree but the entire graph is not a partial 2-tree. The union of the middle component with any other component is also a 2-sum component but not minimal



$G_s = (V_s, E_s)$ of G ; G is *underconstrained* or *independent* and *not rigid* if we have $|E| < 2|V| - 3$ and $|E_s| \leq 2|V_s| - 3$ for all subgraphs G_s . A graph G is *overconstrained* or *dependent* if there is a subgraph $G_s = (V_s, E_s)$ with $|E_s| > 2|V_s| - 3$. G is *well-overconstrained* or *rigid* if there exists a subset of its edges E' such that the graph $G' = (V, E')$ is well-constrained or minimally rigid. A graph is *flexible* if it is not rigid.

5.2 Graphs with Connected, Convex, Linear Polytope 2D Cayley Configuration Space

We are interested in characterizing graphs G that *always admit* a convex, connected or linear polytope Cayley configuration space, i.e., graphs G for which there is a set F of non-edges such that for all distance assignments $\delta(E)$ (resp., intervals $[\delta^l(E), \delta^r(E)]$), the d -dimensional Cayley configuration space $\Phi_F^d(G, \delta)$ (resp., $\Phi_F^d(G, [\delta^l, \delta^r])$) is convex or connected, or is a linear polytope.

Furthermore, given G in this characterized class, we would like to characterize the corresponding sets F of non-edges. These are exactly the parameter choices that yield well-behaved Cayley configuration spaces.

Figure 3 gives an example graph that admits a connected, convex and linear polytope 2D Cayley configuration space on the specified non-edges; and vice versa, in Fig. 5, we provide an example in which the graph does *not* admit such a Cayley configuration space on any non-edge. The graph characterization of Theorem 5.1 can be easily verified for both examples.

Fig. 5 No non-edge f exists such that $\Phi_f^2(G, \delta)$ is always connected




Fig. 6 Base Case 1 of Theorem 5.2





Fig. 7 Base Case 2 of Theorem 5.2: The vertices $u_i, i = 1, \dots, m$, where $m \geq 1$, are the only vertices other than v_1 and v_2 with degree two and they are adjacent to both v_1 and v_2 ; $f = (v_1, v_2)$ is not an edge of the graph



5.2.1 Graphs and Their “Single-Interval” Non-Edges

The next theorem characterizes a graph G along with a non-edge f such that the Cayley configuration space of G on f is always a single interval.


Theorem 5.1 *Given a graph $G = (V, E)$ and a non-edge f , the Cayley configuration space $\Phi_f^2(G, \delta)$ is a single interval for all δ if and only if all the minimal 2-sum components of $G \cup f$ that contain f are partial 2-trees.*

Structure of Proof To prove one (harder) direction of Theorem 5.1, we need the following purely graph-theoretic theorem and the following Lemma 5.5. The other (easy) direction follows from Lemma 5.8 which is, in turn, proven gradually using Lemmas 5.6 and 5.7. □

Theorem 5.2 *Given a graph $G = (V, E)$ and a non-edge f , G can be reduced to the base cases in Figs. 6 or 7 by a sequence of edge contractions (no edge removals) if and only if there exists a minimal 2-sum component of $G \cup f$ containing f that is not a partial 2-tree.*

Remark 5.3 The two base cases Figs. 6 and 7 are based on K_4 . Based on the fact that partial 2-trees do not have K_4 minors and properties of 2-sum, we can prove one direction of Theorem 5.2. For the other direction, the existence of a K_4 minor alone

Fig. 8 G has a connected 2D Cayley configuration space on $f = (v_1, v_2)$ if and only if, for all $1 \leq i \leq k$, the graph G_i has a connected 2D Cayley configuration space on f



is insufficient. We require a special type of pure edge-contraction reduction without edge removals, which additionally preserve selected non-edges, i.e., prevent them from becoming edges or from collapsing to a single vertex.

Proof of Theorem 5.2 (\Rightarrow) We first prove that G cannot be reduced to Figs. 6 or 7 by edge contractions if all the minimal 2-sum components of $G \cup f$ containing f are partial 2-trees. Because partial 2-trees do not have K_4 minors, and since K_4 exists as a minor in both Figs. 6 and 7, we cannot reduce G to either of the two base cases by edge contractions (no edge removals). In fact, in case there exists a 2-sum component $G \cup f$ that does not contain f , our proof will not change since edge contractions either preserve 2-sum relationship or transform a 2-sum to a 1-sum.

(\Leftarrow) We prove the other direction by induction on the number n of vertices of G . The statement is true for the 2 base cases. Assume the statement is true for $|V| \leq n - 1$; we prove it for $|V| = n$. First, we remove v_1 and v_2 to get a set of connected components H_1, \dots, H_k (Fig. 8). We use G_i to denote the subgraph of G which is induced by vertices of H_i together with v_1 and v_2 , where $f = (v_1, v_2)$. Note that each $G_i \cup f$ is a 2-sum component of $G \cup f$. Without loss of generality, we assume $G_1 \cup f$ is one of these 2-sum components of $G \cup f$ but not a partial 2-tree.

Case $k > 1$. If $k > 1$, the number of vertices of $G_1 \cup f$ is $\leq n - 1$. Hence, by the induction hypothesis, we can reduce G_1 to the one of the two base cases by edge contractions. Now we just contract all the edges of H_i when $i > 1$ and the resulting graph falls into the base case in Fig. 7.

Case $k = 1$. In this case, there is only one minimal 2-sum component C containing f , and it is not a partial 2-tree. If C is not $G \cup f$, the number of vertices of C is $\leq n - 1$. By the induction hypothesis, we can reduce C to one of the two base cases. By contracting all the edges of G which are not in C , we can reduce G to one of the base cases. (Note that a 1-sum is a special case of a 2-sum.) If, on the other hand, $C = G \cup f$, then it is a minimal 2-sum component containing f and it is not a partial 2-tree. There are 2 subcases based on l , the maximum number of disjoint paths between v_1 and v_2 , the vertices of f .

[Subcase $l \leq 1$] we can find a vertex, say v_3 , that separates v_1 and v_2 , that is, we can split the graph into two subgraphs G_1 and G_2 such that $v_1 \in G_1, v_2 \in G_2$; G_1 and G_2 share only vertex v_3 , and all the edges of G are either in G_1 or G_2 (refer to Fig. 9 for this case). Since $G \cup f$ is a minimal 2-sum component containing both v_1 and v_2 , both (v_1, v_3) and (v_2, v_3) have to be non-edges. In addition, at least one of $G_1 \cup (v_1, v_3)$ and $G_2 \cup (v_2, v_3)$ is not a partial 2-tree, otherwise $G \cup f$ will also be a partial 2-tree. Without loss of generality, suppose $G_1 \cup (v_1, v_3)$ is not a partial 2-tree. By the induction hypothesis, we can reduce G_1 to one of the two base cases.

Fig. 9 v_3 is an articulation vertex for v_1 and v_2





Fig. 10 Case $k = 1$, subcase $l \geq 2$ in the proof of Theorem 5.2. There are at least two disjoint paths from v_1 to v_2 ; $G \cup f$ is a minimal 2-sum component containing both v_1 and v_2 ; and $G \cup f$ is not a partial 2-tree



By contracting all the edges in G_2 , we can also reduce G to one of the two base cases (v_3 is identified with v_2).

[Subcase $l \geq 2$] Take two disjoint paths $(v_1, t_1, \dots, t_s, v_2)$ and $(v_1, z_1, \dots, z_r, v_2)$. We can contract (t_1, \dots, t_s) to t_1 and (z_1, \dots, z_r) to z_1 (Fig. 10). If we further remove v_1, v_2, t_1 and z_1 , we get new connected components. Then we contract all the edges inside these new connected components such that each of them becomes a single vertex that we denote by q_1, \dots, q_m . Before we contract paths (t_1, \dots, t_s) and (z_1, \dots, z_r) , if we remove v_1 and v_2 , the remaining graph is still connected ($k = 1$), so at least one of q_1, \dots, q_m is connected to both t_1 and z_1 .

Now we contract edges as follows (refer to Fig. 10):

1. If q_i connects to both t_1 and z_1 , we can identify q_i with t_1 by edge contraction.
2. If q_i connects to only v_1 and v_2 (not to t_1 or z_1), we leave it unchanged.
3. If q_i connects to only one of v_1, v_2, t_1 and z_1 , we identify it with the corresponding vertex in v_1, v_2, t_1 and z_1 .
4. If q_i connects to v_1, v_2, t_1 , we can identify q_i with t_1 .
5. If q_i connects to v_1, v_2, z_1 , we can identify q_i with z_1 .

That covers all the cases and completes the proof of the induction step. □

Theorem 5.2 gives us the following independently interesting corollary.

Corollary 5.4 *Given a graph $G = (V, E)$ and a non-edge $f = (v_1, v_2)$, we can reduce G to one of the base cases in Figs. 6 or 7 by a sequence of edge contractions provided the following hold.*

1. G itself is the minimal 2-sum component containing f .
2. For any vertex v_i other than v_1 and v_2 , either $\deg(v_i)$ is 2 and v_i is adjacent to both v_1 and v_2 , or $\deg(v_i)$ is at least three.
3. At least one vertex v_i other than v_1 and v_2 has degree of three or more.

Proof The proof is by contradiction. Suppose (1), (2), (3) hold but G cannot be reduced to either of the two base cases. Then, by Theorem 5.2, all the minimal 2-sum components of the graph $G \cup f$ containing v_1 and v_2 are partial 2-trees. Note that at least one vertex v_i other than v_1 and v_2 has degree of three or more. We consider the 2-sum component C containing v_i . Note also that C has more than 3 vertices, and within C , since (2) holds, there can be no vertices other than v_1 and v_2 that have degree of two or less. By (2), any other vertex of degree two or less would be adjacent to both v_1 and v_2 and would form its own 2-sum components. \square

Now note that, in any 2-tree with more than 3 vertices, there exist two non-adjacent vertices which both have degree of two, so, in any partial 2-tree with more than 3 vertices, there should also exist two non-adjacent vertices which both have degree of at most two. This contradicts the fact that v_1 and v_2 (which are adjacent) are the only vertices in C that have degree of two or less. \square

Using Corollary 5.4, we see that the graph in Fig. 5 can be reduced to one of the two base cases no matter which non-edge we choose.

Lemma 5.5 *In both Figs. 6 and 7, there exists a distance assignment δ s.t. $\Phi_f^2(G, \delta)$ is not connected.*

Proof Follows from reflection across edge (w_1, w_2) . Let $\delta(v_1, w_1)$, $\delta(v_1, w_2)$, $\delta(v_2, w_1)$, $\delta(v_2, w_2)$, and $\delta(w_1, w_2)$ all be 1 and $\delta(v_1, u_i)$ and $\delta(v_2, u_i)$ be 2; we can easily check that possible values of $\delta(f)$ are 0 or $\sqrt{3}$, so $\Phi_f^2(G, \delta)$ is not connected. \square

As noted earlier, Theorem 5.2 and Lemma 5.5 have proved the difficult direction for Theorem 5.1. The following lemmas prove the easy direction.

Lemma 5.6 *If $G = (V, E)$ is the 2-sum of $G_i = (V_i, E_i)$, then, for any δ , (G, δ) has a realization if and only if each (G_i, δ) (δ restricted to the edges in G_i) has a realization.*

Proof of Lemma 5.6 Simply hinge all the 2-sum components' realizations along the 2-sum edges to get a realization of (G, δ) , with any one of two reflection choices across the 2-sum edge. The other direction is immediate. \square

Lemma 5.7 *Take a graph $G = (V, E)$ with 2-sum components $G_i = (V_i, E_i)$, and a non-edge set F that is entirely contained in an arbitrary one of the $G_i = (V_i, E_i)$, say G_1 . Then, for any δ , either $\Phi_F^2(G, \delta) = \Phi_F^2(G_1, \delta)$ if all the $\Phi_F^2(G_i, \delta)$'s are non-empty, i.e., the EDCS (G_i, δ) have at least 1 realization; or otherwise, $\Phi_F^2(G, \delta)$ is empty.*

Proof Directly follows from Lemma 5.6. \square

Lemma 5.8

(a) *If a graph $G = (V, E)$ has a 2-sum component $G' = (V', E')$ that is an under-constrained partial 2-tree, then there exists a nonempty non-edge set F entirely*

- in G' such that for any δ , $\Phi_F^2(G, \delta)$ is a linear polytope. Moreover, there is such a set F such that $\Phi_F^2(G', \delta)$ is generically complete for G' .
- (b) If a graph $G = (V, E)$ is an underconstrained partial 2-tree, then for any non-empty non-edge set F' that preserves $(V, E \cup F')$ as a partial 2-tree, and for all δ , $\Phi_{F'}^2(G, \delta)$ is a linear polytope.

Proof To prove (a), we will construct a nonempty set of non-edges F in G' and show that the projection on F is a linear polytope for all δ .

Note that G' is an underconstrained partial 2-tree, so we can find a nonempty subset of non-edges of G' by adding which we get a 2-tree. We let F be this nonempty set. Note that such an F is a *completion* for G' , i.e., makes G' minimally rigid. Hence, we know that $\Phi_F^2(G', \delta)$ is of full-measure and generically complete, proving the last sentence of the theorem.

To get the linear polytope, note that a 2-tree can be written as the 2-sum of triangles. For example, let $\delta(v_i, v_j)$, $\delta(v_j, v_k)$ and $\delta(v_k, v_i)$ denote the length of the three edges of the triangle $\Delta v_i v_j v_k$, then the triangle inequalities are $\delta(v_i, v_j) \leq \delta(v_i, v_k) + \delta(v_j, v_k)$, $\delta(v_i, v_k) \leq \delta(v_j, v_k) + \delta(v_i, v_j)$ and $\delta(v_j, v_k) \leq \delta(v_i, v_k) + \delta(v_i, v_j)$.

Thus, for all δ , $\Phi_F^2(G', \delta)$ is a linear polytope. Now since F is entirely in G' , Lemma 5.7 applies and for all δ , $\Phi_F^2(G, \delta) = \Phi_F^2(G', \delta)$ or $\Phi_F^2(G', \delta)$ is empty. Thus, for all δ , $\Phi_F^2(G, \delta)$ is also a linear polytope.

To prove (b), for any underconstrained partial 2-tree $G = (V, E)$, we can find a nonempty non-edge set F that makes $(V, E \cup F)$ a 2-tree; and we showed that, for any δ , $\Phi_F^2(G, \delta)$ is a generically complete linear polytope 2D Cayley configuration space. Take any nonempty subset of F' of such an F —these are exactly the subsets of non-edges whose addition would preserve the partial 2-tree property of G . Then $\Phi_{F'}^2(G, \delta)$ is the projection of $\Phi_F^2(G, \delta)$ on F' , and since the latter is a linear polytope, the former is a linear polytope as well. □

Proof of Theorem 5.1 The proof of one (harder) direction follows directly from Theorem 5.2 and the Lemma 5.5. Specifically, to pick a distance assignment δ for G that yields a disconnected Cayley configuration space on f , we set all the contracted edges during the procedure of Theorem 5.2 to 0. The uncontracted edges are now mapped by the reduction to edges of one of the base cases. Lemma 5.5 tells us how to choose those distance values to ensure disconnectedness of the Cayley configuration space on f . The other (easy) direction is immediate from Lemma 5.8. □

Next we show that Theorem 5.1 is tight in that neither of the two straightforward extensions to 3D hold.

Observation 5.9 *There exists a partial 3-tree (resp., 3-realizable graph) $G = (V, E)$ and non-edge f such that $G \cup f$ is not a partial 3-tree (resp., 3-realizable graph), and yet $\Phi_f^3(G, \delta)$ is always connected.*

Proof Refer to and Figs. 11 and 12. □

Fig. 11 For Observation 5.9. G 's 3D Cayley configuration space on non-edge f is one interval, although $G \cup \{f\}$ has a K_5 minor





Fig. 12 For Observation 5.9. Graph G has connected 3D Cayley configuration space on non-edge f ; G does not have K_5 minor or $K_{2,2,2}$ minor; $G \cup f$ has a $K_{2,2,2}$ minor but does not have a K_5 minor; in particular, without contracting edge f in $G \cup f$ we cannot get a $K_{2,2,2}$ minor

5.2.2 Graphs with Generically Complete, Linear Polytope Cayley Configuration Spaces

In Theorem 5.10, we give an exact characterization of the class of graphs G , all of whose corresponding EDCS (G, δ) admit a 2D (*generically complete*), *linear polytope* Cayley configuration space. The theorem also shows that the characterization remains unchanged if the Cayley configuration space is merely required to be convex, and further if it is merely required to be connected.

Theorem 5.10

- (a) For a graph $G = (V, E)$, the following four statements are equivalent:
 1. There exists a non-empty set of non-edges F such that for all δ $\Phi_F^2(G, \delta_E)$ is connected.
 2. There exists a non-empty set of non-edges F such that for all δ , $\Phi_F^2(G, \delta_E)$ is convex.
 3. There exists a non-empty set of non-edges F such that for all δ $\Phi_F^2(G, \delta_E)$ is a linear polytope.
 4. G has a 2-sum component that is an underconstrained partial 2-tree.
- (b) An underconstrained graph G always admits a generically complete linear polytope, connected or convex Cayley configuration space if and only if every underconstrained 2-sum component of G is a partial 2-tree.

Proof For (a), we will prove the cycle $(4) \Rightarrow (3) \Rightarrow (2) \Rightarrow (1) \Rightarrow (4)$. We proved $(4) \Rightarrow (3)$ in Lemma 5.8. A linear polytope is convex, so $(3) \Rightarrow (2)$ follows. Convexity implies connectedness, so $(2) \Rightarrow (1)$ follows. Theorem 5.1 and the proof of Lemma 5.8 proves $(4) \Rightarrow (1)$; therefore, we proved $(1) \Rightarrow (4)$ as well.

For one direction of (b): if every underconstrained 2-sum component of G is an underconstrained partial 2-tree, then by Lemma 5.8, G always admits a generically

complete linear polytope, connected or convex Cayley configuration space. The reverse direction of (b) follows from (a) (1, 2, 3 \Rightarrow 4): If G always admits a generically complete linear polytope, connected or convex Cayley configuration space, G has at least one 2-sum component which is an underconstrained partial 2-tree. \square

5.2.3 Full Characterization of Cayley Parameters that Yield a Linear Polytope 2D Cayley Configuration Space

The following theorem is a refined version of Theorem 5.10.

Theorem 5.11 *Given a graph $G = (V, E)$ and non-empty set of non-edges F , the 2D Cayley configuration space $\Phi_F^2(G, \delta)$ is a linear polytope, connected or convex for all δ if and only if all the minimal 2-sum components of $G \cup F$ containing any subset of F are partial 2-Trees. Furthermore, the Cayley configuration space is generically complete if and only if all the underconstrained minimal 2-sum components of G are partial 2-trees and all the minimal 2-sum components of $G \cup F$ containing F are 2-trees.*

Proof Directly from Theorem 5.1 and the proof of Lemma 5.8. \square


5.2.4 Characterization of EDCS with Distance Intervals

We have characterized graphs G and sets of non-edges F where the Cayley configuration space $\Phi_F^d(G, \delta)$ is connected, convex and a linear polytope for all δ . We can extend the results for distance interval constraints. The next observation and theorem show that while one direction of the above results extends to interval constraints, the other direction fails in the current form. This motivates a modified characterization theorem.

Observation 5.12 *There exists a graph $G = (V, E)$ and non-edge f such that all the minimal 2-sum components of $G \cup f$ containing f are partial 2-trees, and yet we can find an interval distance constraint $[\delta^l, \delta^r]$ such that $\Phi_f^2(G, [\delta^l, \delta^r])$ is not a linear polytope, convex, or connected.*

Proof Refer to Fig. 13. Denote the graph shown in Fig. 13(a) by G and the graph shown in Fig. 13(b) by G' . Denote the non-edge (v_1, v_2) by f in both Figs. 13(a) and 13(b). Clearly, $G \cup f$ has 3 minimal 2-sum components: the subgraph induced by (v_1, v_2, v_9, v_{10}) , the subgraph induced by (v_1, v_2, v_3, v_4) and the subgraph induced by $(v_1, v_2, v_3, v_4, v_5, v_6, v_7, v_8)$. The first two minimal 2-sum components both contain f and both are partial 2-trees. The third one (the subgraph induced by $(v_3, v_4, v_5, v_6, v_7, v_8)$) is not a partial 2-tree but does not contain f . By Theorem 5.1, for all distance δ , the Cayley configuration space $\Phi_f^2(G, \delta)$ is connected. Similarly, $G' \cup f$ has 2 minimal 2-sum components: the subgraph induced by (v_1, v_2, v_9, v_{10}) , and the subgraph induced by $(v_1, v_2, v_3, v_4, v_5, v_6, v_7, v_8)$. Both two minimal 2-sum components contain f . The first one is not a partial 2-tree and the second one (the subgraph induced by $(v_1, v_2, v_3, v_4, v_5, v_6, v_7, v_8)$) is not a partial 2-tree. By Theorem 5.1, we can find a distance δ' such that the Cayley configuration space $\Phi_f^2(G, \delta')$

Fig. 13 For Observation 5.12. All the minimal 2-sum components of $G \cup f$ containing f are partial 2-trees, but $\Phi_f^2(G, [\delta^l(E), \delta^r(E)])$ is not always connected



is not connected. If we allow an interval distance constraint on edge (v_3, v_4) in G , we would have an equivalent distance constraint system as G' . \square

Now we give a characterization theorem for distance interval constraints.

Theorem 5.13 *Given a connected graph $G = (V, E)$ and a nonempty set of non-edges F , the 2D Cayley configuration space $\Phi_F^2(G, [\delta^l, \delta^r])$ is a linear polytope, convex or connected for all $[\delta^l, \delta^r]$ if and only if every minimal 2-sum component of a specific subdivision G'_F of $G \cup F$ that contains any subset of F is a partial 2-tree. We construct the subdivision G'_F of $G \cup F$ by replacing each edge (u_1, u_2) by a path of length 2 by introducing a new vertex u .*

Proof For a given interval constraint $[\delta^l(u_1, u_2), \delta^r(u_1, u_2)]$, we can choose $\delta(u_1, u) = \frac{|\delta^r(u_1, u_2) - \delta^l(u_1, u_2)|}{2}$ and $\delta(u, u_2) = \frac{\delta^r(u_1, u_2) + \delta^l(u_1, u_2)}{2}$ such that the distance constraints on the subdivision, namely $\delta(u_1, u)$ and $\delta(u, u_2)$, together impose the original distance interval constraint on the edge (u_1, u_2) . Then, by Theorem 5.11, $\Phi_F^2(G', \delta)$ is a linear polytope, connected or convex for all δ if and only if all the minimal 2-sum components of G'_F containing a subset of F are partial 2-trees. \square

5.3 Graphs with Universally Inherent, Convex, Squared Cayley Configuration Spaces

5.3.1 d -Realizability Implies Convex, Squared Cayley Configuration Spaces

We will prove in Theorem 5.15 that d -realizable graphs admit universally inherent, connected (resp., convex) d -dimensional (resp., squared) Cayley configuration spaces. Before that, we first show in Lemma 5.14 that convexity of squared Cayley configuration spaces implies connectedness of the Cayley configuration space.

Note All the results in the section apply to both distance constraints δ as well as distance interval constraints $[\delta^l, \delta^r]$. In order to avoid writing “for all δ (resp., $[\delta^l, \delta^r]$) the (resp., squared) Cayley configuration space $\Phi_F^d(G, \delta)$ (resp., $\Phi_F^2([\delta^l, \delta^r])$) is connected (resp., convex)” we instead just say “ G always admits a connected (resp., convex) (resp., squared) Cayley configuration space.”

Lemma 5.14 *If a graph always admits universally inherent, convex, d -dimensional squared Cayley configuration spaces, then it also always admits universally inherent, connected, d -dimensional configuration spaces.*

Proof For a non-edge set $F = \{f_1, f_2, \dots\}$ of G , denote by $(\Phi_F^d)^2(G, \delta)$ the Cayley configuration space $\{((\delta^*)^2(f_1), (\delta^*)^2(f_2), \dots) : (G \cup F, \delta, \delta^*(F)) \text{ has a realization in } \mathbb{R}^d\}$. The map $(\cdot)^2 : \Phi_F^d(G, \delta) \rightarrow (\Phi_F^d)^2(G, \delta)$ is continuous, and the inverse map is well-defined and continuous over the positive reals. Now the convexity of $(\Phi_F^d)^2(G, \delta)$ implies its connectedness that, by the above-mentioned continuity, implies the connectedness of $\Phi_F^d(G, \delta)$. \square

Now we are ready to give the theorem and the proof.

Theorem 5.15 *If a graph is d -realizable, it admits universally inherent, connected (resp., convex), d -dimensional (resp., squared) Cayley configuration spaces.*

Proof By Lemma 5.14, we only need to prove a d -realizable graph admits universally inherent, convex, squared Cayley configuration spaces.

An $n \times n$ matrix M is a *Euclidean square distance matrix (EDM)* if $\exists p_1, \dots, p_n \in \mathbb{R}^d$ for some d such that $\|p_i - p_j\|^2 = M(i, j)$. A classical result [33] that follows from positive semidefiniteness of Gram matrices is that the set of all EDMs is a convex cone (note that d , and hence the rank of these matrices, is not fixed). The projection of this cone on any set $E \cup F$ of pairs (i, j) is also convex. By the definition of d -realizability of a graph $H = (V, E \cup F)$ with $|V| = n$, this projection is exactly the set of all squared distance assignments $(\delta^*)^2$ to the pairs in $E \cup F$ for which (H, δ^*) has a realization in \mathbb{R}^d . We denote this by $(\Phi_{E \cup F}^d)^2((V, \phi), \phi)$. Since convexity is preserved by both sections and projections, the section of this projection $(\Phi_{E \cup F}^d)^2((V, \phi), \phi)$ —obtained by fixing δ^* to be δ over E —is also convex, for all δ . This section is exactly the Cayley configuration space $(\Phi_F^d)^2(G, \delta)$ of the graph $G := (V, E)$ over its non-edge set F . Hence, this is a convex squared Cayley configuration space. Note that this holds for any partition $E \cup F$ of the edge set of the d -realizable graph H . Hence, H always admits universally inherent, convex, squared Cayley configuration spaces. \square

Theorem 5.15 gives one direction for all dimensions. We conjecture in Sect. 6.1.1 the reverse direction is also true. In the next section, we will prove that the reverse direction is true for $d \leq 3$.

5.3.2 When Does Universally Inherent, Connected Configuration Space Imply d -Realizability?

Theorem 5.16 *Let $d \leq 3$. Then the following are equivalent for a graph H .*

1. H is d -realizable.
2. H always admits universally inherent, connected, d -dimensional Cayley configuration spaces.
3. H always admits universally inherent, convex, d -dimensional squared Cayley configuration spaces.

Remark 5.17 The above theorem is a weak statement for $d \leq 2$. For example, since the class of 2-realizable graphs is exactly the partial 2-trees, Theorem 5.11 shows

that if a graph H is not 2-realizable, then it has a minimal 2-sum component that is not 2-realizable (not a partial 2-tree). And this component does not admit *any* inherent Cayley configuration space, let alone universally inherent ones. In other words, in this non-2-realizable minimal 2-sum component H_C , on a vertex set V_C for every partition of edges into $E_C \cup F_C$, the Cayley configuration space $\Phi_{F_C}^2(G_C, \delta)$ of graph $G_C = (V_C, E_C)$ is disconnected. For $d = 3$, on the other hand, no such strong statement holds as shown in the counterexample of Observation 5.9. To show the above theorem, we merely show that if a graph H is not 3-realizable, then *there exists* a partition of the edges of H into $E \cup F$, such that the configuration space $\Phi_F^3(G, \delta)$ of the graph $G = (V, E)$ is disconnected.

Structure of Proof The proof of the above theorem requires several lemmas. The idea of the proof is as follows. Lemma 5.14 proves (3) \Rightarrow (2), and Theorem 5.15 proves (1) \Rightarrow (3) for all dimensions d . Based on the above remark, we restrict ourselves to $d = 3$, and just prove (2) \Rightarrow (1). \square

For any non-3-realizable graph $H = (V, E')$, we find a partition of E' as $E \cup F$ where G is the graph (V, E) and find a distance assignment δ such that the Cayley configuration space $\Phi_F^3(G, \delta)$ is disconnected. Here $G := (V, E)$ and f is a single non-edge of G , so this is a 1-parameter Cayley configuration space, and we show that it has 2 isolated points. To do this, we start from the following theorem.

Theorem 5.18 (Connelly and Slougher [3, 4]) *A graph is 3-realizable if and only if it does not have K_5 or $K_{2,2,2}$ as minors.*

Given the existence of one of these 2 minors, as in Theorem 5.2, we show how to pick from H the graph G and its non-edges F such that by a restricted reduction that uses only edge contractions (no edge removals) and preserving the non-edges F , we can reduce the graph G to a K_5 or $K_{2,2,2}$ that is missing exactly one edge, namely f , on to which all the non-edges in F have been mapped by the reduction. Finally, we obtain using Lemma 5.21 the distance assignment δ for G by setting the distances for the contracted edges to 0; and, similar to Lemma 5.5 we pick distance assignments for the un-contracted edges in such a way that the two base cases do not have connected 3D Cayley configuration spaces on f .

We will use the following simple fact.

Fact 5.19 *If a graph $G = (V, E)$ has a complete graph K_m as a minor, then G can be reduced to K_m by edge contraction alone (and removal of isolated vertices if necessary), without edge removals.*

We also need the following independently interesting theorem which strengthens the forbidden-minor theorem of Theorem 5.18.

Theorem 5.20 *If a graph is not 3-realizable, it can be reduced to K_5 or $K_{2,2,2}$ by edge contractions alone (no edge removals).*

Proof If a graph is not 3-realizable, it has a K_5 or $K_{2,2,2}$ minor by Theorem 5.18. If a graph has a K_5 minor, using Fact 5.19, it can be reduced to K_5 by edge contractions alone, so we only need to prove the case that G has a $K_{2,2,2}$ as minor. If G has a $K_{2,2,2}$ as minor, by definition of minor, we can get a $K_{2,2,2}$ by first contracting some edges, then removing some edges, and finally removing some isolated vertices. Follow this reduction path, but stop after edge contractions, and denote the new graph by G' . Denote the $K_{2,2,2}$ subgraph of G' by M .

Now we will show we can either get M or a K_5 by edge-contractions and removing some isolated vertices if necessary. The strategy is straightforward: successively contract one edge at a time, whose two vertices are not both in M until we cannot continue. After removing any possible isolated vertices, the remaining graph has exactly 6 vertices (the same as M) and we denoted this graph by M' . We know that $K_{2,2,2}$ is a subgraph of M' . Now we use a simple observation that by adding one or more edges to $K_{2,2,2}$, we can get a K_5 by edge contractions alone, thus G can be reduced to K_5 or $K_{2,2,2}$ by edge contractions alone. \square

For proving the next lemma, we give an appropriate distance assignment δ to K_5 and $K_{2,2,2}$ to show they do not admit universally inherent, connected Cayley configuration spaces.


Lemma 5.21 *Neither $K_{2,2,2}$ nor K_5 always admits universally inherent, connected, 3D Cayley configuration spaces.*

Proof To prove the Lemma for K_5 (resp., $K_{2,2,2}$), we take one edge as f and prove that we can find a distance assignment δ for graph $K_5 \setminus f$ (resp., $K_{2,2,2} \setminus f$) such that $\Phi_f^3(K_5 \setminus f, \delta)$ (resp., $\Phi_f^3(K_{2,2,2} \setminus f, \delta)$) has more than one interval.

For the K_5 case, we take f to be (v_1, v_2) and assign the same distances to all the edges in $K_5 \setminus f$. If we fix tetrahedron (v_1, v_3, v_4, v_5) , then v_2 can either be coincident with v_1 or is the reflection of v_1 about plane (v_3, v_4, v_5) . Since v_1 is not in the plane (v_3, v_4, v_5) , so in the latter case, $\delta^*(f)$ is not zero. In the former case, $\delta^*(f)$ is zero. These two values are all the possible values of $\delta^*(f)$, so we have proved that $\Phi_f^3(K_5 \setminus f, \delta)$ is not connected.

For the $K_{2,2,2}$ case (Fig. 14), we choose edge (v_5, v_6) as f . We choose a distance assignment δ so that the following conditions are satisfied: $\delta(v_1, v_2) = \delta(v_2, v_3) = \delta(v_1, v_3) = \delta(v_1, v_4) = \delta(v_3, v_4) = \delta(v_2, v_5) = \delta(v_3, v_5) = \delta(v_4, v_5)$, $\delta(v_1, v_4) + \delta(v_4, v_6) = \delta(v_1, v_6)$, $\delta(v_4, v_6) > 0$, and $\delta(v_2, v_6)$ will let $\angle(v_2, v_1, v_6) = \frac{\pi}{3}$.

Fig. 14 (a) $K_{2,2,2}$ with one edge f removed; (b) see proof of Lemma 5.21: A distance assignment δ to $K_{2,2,2} \setminus f$ such that the 3D Cayley configuration space on f is not connected



Because $\delta(v_1, v_4) + \delta(v_4, v_6) = \delta(v_1, v_6)$, v_1, v_4 and v_6 are collinear and the four vertices v_1, v_2, v_4 and v_6 are coplanar. Because $\angle(v_2, v_1, v_6) = \frac{\pi}{3}$ and $\delta(v_1, v_2) = \delta(v_1, v_4)$, $\delta(v_2, v_4)$ will be equal to $\delta(v_1, v_2)$ and $\delta(v_1, v_4)$. Thus, the location of v_3 will be uniquely determined by the normal tetrahedron (v_1, v_2, v_3, v_4) if we assume that, without loss of generality, v_3 is above plane (v_1, v_2, v_4, v_6) . Now v_5 has two possible locations, either v_1 or the reflection of v_1 about plane (v_2, v_3, v_4) . We denote the former location as v_5 while the latter as v'_5 . For the case when v_5 is coincident with v_1 , $\delta(v_5, v_6)$ will be $\delta(v_5, v_6) = \delta(v_1, v_6) = \delta(v_1, v_4) + \delta(v_4, v_6)$. For the other case, by the triangle inequality in $\Delta(v_4, v'_5, v_6)$, we have $\delta(v'_5, v_6) \leq \delta(v'_5, v_4) + \delta(v_4, v_6) = \delta(v_1, v_4) + \delta(v_4, v_6) = \delta(v_1, v_6) = \delta(v_5, v_6)$. Now (v_1, v_2, v_3, v_4) and (v_2, v_3, v_4, v'_5) are normal tetrahedrons and v'_5 is different from v_1 in this case, so v'_5 is not in the plane (v_1, v_2, v_4) . Therefore, v'_5 cannot be collinear with v_4 and v_6 , so $\delta(v'_5, v_6) > \delta(v_5, v_6)$. Thus, $\delta(f)$ has two disconnected values, and we have proved that projection $\Phi_f^3(G, \delta)$ is not connected. \square

Now we are ready to prove Theorem 5.16.

Proof of Theorem 5.16 By Theorem 5.15, a 3-realizable graph H admits universally inherent connected, convex, 3D squared Cayley configuration spaces, so we only need to prove the reverse direction. If H is not 3-realizable, by Theorem 5.20, we can get a K_5 or $K_{2,2,2}$ by edge contractions alone.

As mentioned before, find a partition of the edge set of H into $F \cup E$ that defines a subgraph $G = (V, E)$ and a non-edge set F for G . Then find a distance assignment δ for E such that $\Phi_F^3(G, \delta)$ is disconnected.

Since H can be reduced to K_5 or $K_{2,2,2}$ by edge contractions alone, we pick an edge from the corresponding minor and denote it f . We choose F to be all the edges that were identified with f by the reduction. For the distance assignment δ to the edges of G , we will use a similar method as in the proof of Theorem 5.1. Lemma 5.21 gives a distance assignment for $K_{2,2,2}$ and K_5 that ensures that the Cayley configuration space on the edge f is disconnected. Set the distances of each uncontracted edge e during the reduction of G to the distance assignment of the edge in K_5 or $K_{2,2,2}$ that e was identified with. Set the distances of all the contracted edges of G to be 0. This ensures that $\Phi_F^3(G, \delta)$ is disconnected. \square

5.4 Efficiently Recognizing Graphs with Connected and Convex Cayley Configuration Spaces

Theorems 5.10 and 5.16 characterize when a given graph always admits (universally inherent) connected, convex, (generically complete), linear polytope 2D and 3D (squared) Cayley configuration spaces. Based on these theorems, we give efficient algorithms to recognize these properties.

Theorem 5.22 *Given a graph G , both for the case of distance constraints as well as for distance interval constraints, there are linear time algorithms to recognize:*

1. *Whether there is a non-empty set of non-edges F such that G always admits a connected (convex, linear polytope) 2D Cayley configuration space on F .*

2. Under the assumption that G is not over-constrained, whether there is a set of non-edges F such that G always admits a connected (convex, linear polytope) generically complete 2D Cayley configuration space on F .
3. Whether G always admits universally inherent, connected (resp., convex), d -dimensional (resp., squared) Cayley configuration spaces, for $d \leq 3$.

Proof For (1), first note that a linear time algorithm decomposes the input graph to 2-sum components. By Theorem 5.10, there is a non-empty set of non-edges F such that G always admits a connected (convex, linear polytope) 2D Cayley configuration space on F , for the case of distance constraints if and only if G has a 2-sum component that is an underconstrained partial 2-tree. Thus, we only need to check whether there exists a 2-sum component that is an underconstrained partial 2-tree. Checking if a graph is a partial 2-tree can be done in linear time. Since partial 2-trees are always independent, i.e., they cannot be overconstrained, they are underconstrained exactly if the number of edges is at most twice the number of vertices minus 3 (the Laman count). Thus, the entire check can be done in linear time.

Under the assumption that G does not have overconstrained subgraphs, the algorithm for (2) is essentially the same as (1). By Theorem 5.1, check whether all underconstrained 2-sum components are partial 2-trees. Find any set of non-edges F that complete them into 2-trees, which are automatically well-constrained or minimally rigid. This guarantees a generically complete Cayley configuration space on F . The assumption that G has no overconstrained subgraphs is necessary for the above algorithm to work. For example, suppose the input graph G is a 2-sum of two graphs G_1 and G_2 where G_1 is an underconstrained partial 2-tree while G_2 is not a partial 2-tree. There is a complete set of non-edges F for G_1 such that G_1 always admits a generically complete, connected, convex and linear polytope 2D Cayley configuration space. However, in order to check whether F yields a generically complete Cayley configuration space for all of G , we have to ensure that G is rigid. If no overconstrained subgraphs exist, then this can be determined by a Laman count for G_2 in linear time. Otherwise, we would need to detect the presence of overconstrained subgraphs. In general, $O(|E||V|)$ network flow based algorithms [16, 34], and even more efficient pebble game algorithms [23, 27] exist for this check.

For (1) and (2) in the case of distance interval constraints, Theorem 5.13 permits us to perform the same checks and in linear time on a subdivision G' of G that is at most 2 times the size of the edge set of G .

For (3), by Theorem 5.16, which applies to both distance and distance interval constraints, we only need to check whether a graph is 2-realizable (partial 2-tree) and 3-realizable. The former is straightforward. Based on [3], a linear algorithm is proposed in [36] to check whether a graph is 3-realizable. \square

5.5 Practical Use: Sampling and Realization Complexities

Theorem 5.22 gives us the algorithms to find a set of non-edges that ensure (generically complete), connected, convex, linear polytope Cayley configuration spaces. They prove Points (1), (2), (3) of the Theorem below. For practical use, we need to further get the exact description of $\Phi_F^d(G, \delta)$ (resp., $\Phi_F^d(G, [\delta^l, \delta^r])$) (sampling complexity) and also a Cartesian realization of a given parametrized configuration in $\Phi_F^d(G, \delta)$ (resp., $(G, [\delta^l, \delta^r])$) (realization complexity).

Theorem 5.23 *Given an EDCS (G, δ) or $(G, [\delta^l, \delta^r])$, where the corresponding graphs G are recognized by the algorithms of Theorem 5.22, there are linear time algorithms that:*

1. *Output a set of non-edges F such that G always admits a (generically complete) connected (convex, linear polytope) 2D Cayley configuration space on F .*
2. *Under assumption G is not overconstrained, output a set of non-edges F such that G always admits a connected (convex, linear polytope) generically complete 2D Cayley configuration space on F .*
3. *Output a maximal set of non-edges F such that the graph $H := G \cup F$ always admits universally inherent, connected 2D or 3D Cayley configuration spaces.*
4. *For F in all of the above items, output a corresponding description of a linear polytope 2D Cayley configuration space $\Phi_F^2(G, \delta)$ (resp., $\Phi_F^2(G, [\delta^l, \delta^r])$) provided it is known to be nonempty; additionally, if one realization p is known, then given any input element in the 2D Cayley configuration space, output its corresponding realization in \mathbb{R}^2 .*

Given an EDCS (G, δ) or $(G, [\delta^l, \delta^r])$, there is a polynomial time algorithm that

5. *For the F in Item 3, outputs one configuration in the configuration space $\Phi_F^3(G, \delta)$ (resp., $\Phi_F^3(G, [\delta^l, \delta^r])$); additionally, when given an input element of this 3D Cayley configuration space, outputs a realization \mathbb{R}^3 .*

Proof As pointed out earlier, the proof of Theorem 5.22 also proves (1), (2). For (3), for 2D, we simply output a set of non-edges F so that $G \cup F$ is a complete 2-tree. For 3D, using a characterization theorem of 3-realizable graphs G , by [3, 4], and an algorithm of [36], we can output a maximal set of edges F such that $G \cup F$ remains 3-realizable. This gives the desired graph H by Theorem 5.15.

For (4), for the case of distance constraints alone, the proof of Lemma 5.8 gives a linear time algorithm to find the description of the linear polytope 2D configuration space on F as a system of linear inequalities. For the case of distance interval constraints, the same algorithm can be applied to the subdivision graph G' given by Theorem 5.13 to get the description. The assumption that the Cayley configuration space is non-empty is required because by Lemma 5.7, the Cayley configuration space of G would be empty if any of EDCS corresponding to a non-partial-2-tree 2-sum components of G has no realization; and determining if a general EDCS has a realization is an NP-hard problem.

For the second statement in (4), take G_C (resp., G'_C) to be any of the 2-sum components of G (resp., G') that contain a subset of $F_C \subseteq F$. When an element, $\delta^*(F_C)$ of a Cayley configuration space of such an EDCS is given, if $G_C \cup F_C$ (resp., $G'_C \cup F_C$) are complete 2-trees, the corresponding EDCS $(G_C \cup F_C, \delta, \delta^*(F_C))$ (resp., $(G'_C \cup F_C, \delta, \delta^*(F_C))$) are Quadratically or Radically realizable, i.e., they can be realized by a straightforward ruler-and-compass construction involving a sequence of solutions of univariate quadratics (intersection of circles). On the other hand, if $G_C \cup F_C$ (resp., $G'_C \cup F_C$) are not complete 2-trees, they can be made complete 2-trees by additional edges D_C , for which δ can be extended consistently, i.e., in such a way that $(G_C \cup D_C \cup F_C, \delta, \delta(D_C), \delta^*(F_C))$ has a realization whenever

$(G_C \cup F_C, \delta, \delta^*(F_C))$ has a realization. The realizations of the remaining 2-sum components (that do not contain any subset of F), in particular those that may not be partial 2-trees, can be read off from the one given realization p . As shown in Lemma 5.6, realizations of the different 2-sum components do not interfere with each other, i.e., the realizations of the various 2-sum components can be hinged together along the 2-sum edges to get the realization of the entire graph. Again, the realization p is required because finding a realization for the EDCS of a 2-sum component that is a general graph (not necessarily a partial 2-tree) is a hard problem (the decision version is NP-hard [32]).

For (5), simply using the definition of 3-realizable graphs, we can obtain one point $\delta^*(F)$ of the Cayley configuration space of G on the given F as follows. Use the positive semidefiniteness of the Euclidean distance matrix cone to complete the distances or distance intervals (for the edges of G as given by δ) into a full Euclidean distance matrix for some dimension d . This d could be much larger than 3. Since we know that the given F satisfies the property that $H = G \cup F$ is also 3-realizable, by the definition of 3-realizability, if we pick the distances for $\delta^*(F)$ from the completed matrix, then it would be guaranteed to be a point in the Cayley configuration space of G on F .

Now given such a $\delta^*(F)$ for a such a maximal set F , using a characterization of [3, 4] of 3-realizable graphs as 2-sums and 3-sums of a small number of special types of graphs, the paper of [36] gives an algorithm to get one realization of $(G \cup F, \delta, \delta^*(F))$ in \mathbb{R}^3 . \square

6 Conclusions and Conjectures

Our results give a practically meaningful, and mathematically robust definition of exact and efficient Cayley configuration spaces of underconstrained Euclidean Distance Constraint Systems (equalities and inequalities), based on various efficiency factors including complexity of sampling and realization. We have taken the first step in a systematic and graded program sketched in [12] and laid out in [11]—for the combinatorial characterizations of efficient Cayley configuration spaces. In particular, the results presented here characterize graphs and their Cayley parameters that yield (squared) Cayley configuration spaces that are connected, convex, linear polytopes, and efficient algorithms for sampling realizations.


6.1 Theoretical Directions and Conjectures

6.1.1 *Is d -Realizability Equivalent to Universally Inherent, Connected d -Dimensional Cayley Configuration Spaces?*

Our first conjecture is the reverse direction of Theorem 5.15.

Conjecture 6.1 *For any dimension d , a graph is d -realizable if and only if it always admits universally inherent, connected (resp., convex) d -dimensional squared Cayley configuration spaces.*

Fig. 15 A generically globally rigid graph in 2D



We can try to leverage results [9, 17] about connected components of the d -dimensional realization spaces of a graph based on its higher, d' -dimensional realization spaces, where $d' > d$.

6.1.2 The Roles of Genericity and Independence

Our results characterize (for the case of distance equalities and frameworks) Cayley configuration space properties that hold for *all* distance assignments δ , and are hence incorrect if we require the properties to hold only in generic situations. For example, Fig. 15 shows a generically globally rigid graph G : The generic 2D Cayley configuration space of this graph on the non-edge f is a single point. However, by our Theorem 5.1, since a minimal 2-sum component containing f is not 2-realizable, the Cayley configuration space on f is disconnected. The apparent discrepancy arises because the proof of the Lemma 5.2 uses non-generic specializations of the edge distances δ in the process of reduction to a minor which then shows that the Cayley configuration space on f is disconnected *for that* δ , which is sufficient to prove the statement of Lemma 5.2 and Theorem 5.1. We believe, however, that both these results still hold under a genericity assumption, *provided* the graph G is not overconstrained (independent, in rigidity terminology) and the non-edge f is not globally-linked.

Conjecture 6.2 *Given a graph G that is not generically over-constrained and a non-edge f , the Cayley configuration space $\Phi_f^2(G, \delta)$ is a single interval for all generic δ (i.e., for δ that admit a 2D generic realization of (G, δ)) if and only if all the minimal 2-sum components containing f are 2-realizable and (partial 2-trees).*

Proving this would formally establish the connections between our work and the work on generic global rigidity and generic globally linked pairs, including results of [9, 17, 20, 21] and [22].

6.1.3 Sampling Complexity for 3D

Note that, while we know from Theorem 5.16 that 3-realizable graphs always admit universally inherent, convex, 3D squared Cayley configuration spaces, we do not yet know an efficient algorithm to determine their description. This is necessary to determine the sampling complexity. This is in stark contrast to the linear time algorithm for obtaining such descriptions in the case of 2D (Theorem 5.11).

One straightforward algorithm for obtaining the description of the Cayley configuration space $\Phi_F^d(G, \delta)$ as a semi-algebraic set is to start with the Cayley–Menger [6, 8, 29] determinantal equalities and inequalities for Euclidean distance matrices in d -dimensions. These are polynomial relationships between all the $\binom{|V|}{2}$ Cayley parameters, including those in E , F and those in $\overline{E \cup F}$. Eliminating those in $\overline{E \cup F}$, and specifying the values for the parameters in E to be $\delta(E)$ leaves polynomial equalities and inequalities in the Cayley parameters in F . This yields the required description of the Cayley configuration space $\Phi_F^d(G, \delta)$ as a semi-algebraic set.

Viewed in this manner, it appears remarkable that in 2D, for the graphs G and non-edge sets F satisfying the conditions of Theorem 5.11, this above-described elimination leaves only the triangle inequalities relating the Cayley parameters in E and F (note that these were part of the original Cayley–Menger set of inequalities), and hence we get a linear polytope description of $\Phi_F^2(G, \delta)$. Note, however, that *we did not use such an elimination for our proof!* Our proof that these triangle inequalities give a description of $\Phi_F^2(G, \delta)$ was through a more direct route: We determined for what configuration $\delta^*(F)$ for the Cayley parameters in F we could construct a 2D realization for the augmented EDCS $\Phi_F^2(G \cup F, \delta(E), \delta(F^*))$.

We would like to show a similar result in 3D either by using elimination or using Euclidean distance matrix completion for fixed rank [1, 5] by a more direct route of determining when 3D realizations can be constructed.

Conjecture 6.3 *Let H be a 3-realizable graph on vertex set V . Take any partition of the edge set of H into $E \cup F$, and consider the graph $G = (V, E)$ and any EDCS (G, δ) . Then there is an $O(|V|^2)$ time algorithm to write down the description of the Cayley configuration space $\Phi_F^2(G, \delta)$ as a semi-algebraic set of low degree (say, no more than 4).*

In this context, note that there is also room to improve the sampling complexity and obtain more refined descriptions even for 2D Cayley configuration spaces of graphs satisfying the conditions of Theorem 5.11. For this purpose, it makes sense to study complete 2-trees H on a set of vertices V . In particular, after partitioning the edges of H into $E \cup F$ and eliminating all except the set of Cayley parameters in E we obtain a system of linear inequalities in those parameters, which we call the *2D admissible distance polytope* of $G = (V, E)$, i.e., it specifies the distance assignments δ for which (G, δ) has a 2D realization. Note that this is the projection of the 2D admissible distance polytope for H , onto the parameters in E . Furthermore, the Cayley configuration space $\Phi_F^2(G, \delta)$ is a *section or a fiber* of this polytope. In [7], we study the detailed topology of these polytopes, towards extending similar studies of polytopes corresponding to the simpler class of polygonal linkages parametrized by Cayley parameters, see, for example, [14].

6.1.4 Characterizing All Parameters that Ensure Connected, Cayley Squared Configuration Spaces for 3D

It would be desirable to extend Theorem 5.1 (and hence Theorems 5.10 and 5.11) to a characterization of parameter choices for 3D connected Cayley configuration

spaces. As pointed out in Observation 5.9, the obvious analog of this result fails in 3D. Partial results supporting the following conjecture have been reported in [11]. Note that partial 3-trees are a large proper subclass of 3-realizable graphs.

Conjecture 6.4 *Given graph G that is a partial 3-tree and non-edge f*

- *If $G \cup f$ has no K_5 or $K_{2,2,2}$ minor, then G has a connected 3D Cayley configuration space on f .*
- *If $G \cup f$ has a K_5 or $K_{2,2,2}$ minor then G has a connected 3D Cayley configuration space on f if and only if the 2 vertices of f must be identified in order to get a K_5 or $K_{2,2,2}$ minor in G .*

6.1.5 1-Dof 2D Mechanisms and 2D Cayley Configuration Spaces with 2 Connected Components

There are 2 possible directions to move beyond the connected, convex, linear-polytope 2D configuration spaces.

The first is to study the Cayley configuration spaces of 1-dof graphs that *just fail* to satisfy the requirements of Theorem 5.2. A first step in this direction has been taken by [13], see also [11], and future research suggestions along that direction has been presented there.

The second direction is to study which graphs always admit Cayley configuration spaces with 2 or fewer connected components? Partial results in this direction have been reported in [11], specifically towards characterizing graphs G and non-edges f such that the 2D Cayley configuration space of G on f has no more than 2 intervals.

6.1.6 Stronger Results for Special Distance Assignments

Our characterizations apply to all distance assignments δ or distance intervals $[\delta^l, \delta^r]$. This places strong restrictions on the characterized classes of graphs G with efficient Cayley configuration spaces. However, if one assumes special distance assignments, well-behaved Cayley configuration spaces may exist for much larger classes of graphs. For example, consider the 2D 2-direction grid which is not a partial 2-tree. However, under the restriction of unit-distance edges, [11] shows that such an EDCS has a Cayley configuration space that is a convex polytope, in fact, a *rectilinear box*. Moreover, for such an EDCS, [11] gives a complete characterization of Cayley parameters (non-edges) on which the Cayley configuration space is convex. It would be desirable to obtain similar results for 3D, due to potential applications in Sect. 6.2 below.

6.2 Practical Implications for Mechanical CAD and Computational Chemistry

6.2.1 Realization Space and Motion Exploration for CAD Mechanisms

Representation of realization spaces of underconstrained systems is a long-standing issue in the development of constraint solvers underlying mechanical CAD systems. Our characterizations and algorithms for obtaining efficient configuration spaces

and sampling realizations would be useful to incorporate into commercial constraint solvers. However, they need to be implemented in combination with other practical user- or designer-driven functionalities.

In addition to sampling realizations, our program of comprehensive study of possible parameter choices is useful for *geometrically meaningful exploration* of the motions of underconstrained systems, or mechanisms. Current (infinitesimal) motion representation reduces to a basis choice for the motion space obtained from the rigidity matrix for a particular realization p . Obvious basis choices lead to an erratic exploration of the motion space, often amplifying spurious “globally” coupled motions or “allosteric” effects instead of systematically first exploring “locally” coupled motions to the extreme or boundary configurations.

Intuitively, our theorems give choices of Cayley parameters that ensure convex or otherwise well-behaved Cayley configuration spaces, and hence each standard method of walking these spaces invests a geometric meaning to the corresponding systematic motion exploration, especially since Cayley parameters are particularly suited to representing internal motion. Furthermore, any standard method of walking a convex Cayley configuration space would fully explore “locally” coupled motions before systematically progressing to less local ones. It would be desirable to state and prove a formal statement to this effect.

6.2.2 Sampling Realization Space for Helix Packing

Helix packing is a well-studied computational chemistry problem since helices are prevalent in many biomolecules. One part of the problem can be viewed as sampling the realization space of an EDCS (including distance inequalities) that correspond to steric or collision avoidance and other constraints. Current methods use inexact representations of realization spaces and “generate and test” algorithms for sampling them. As a result, they are inefficient and moreover lack any guarantee that all boundary and extreme realizations have been explored.

Partial results [11] (see Sect. 6.1.4) give parameter choices that ensure connected and convex, 3D squared Cayley configuration spaces *even for non 3-realizable graphs*. By the characterization proved here, these are clearly not universally inherent. Based on these Cayley parameter choices, an algorithm for exact description and efficient, systematic sampling of the helix packing configuration spaces (especially the boundaries and extreme points) is given in [35].

6.2.3 Limit Properties and Sampling of Realization Spaces for Zeolite Sequences

Zeolites are naturally occurring materials whose structure can be described as “corner sharing regular tetrahedra.” The corresponding EDCS sequences involve only distance equalities and are obtained from progressively larger pieces of an infinite lattice structure with uniform boundary conditions. The goal is to sample the realization spaces of these Zeolite EDCS sequences, with a view to understanding their properties in the limit.

Since the Zeolite EDCS have special distance assignments, as pointed out in Sect. 6.1.6, they have efficient Cayley configuration spaces even if the corresponding

graphs may not satisfy the requirements of the characterization theorems presented in this manuscript. Using similar results such as a characterization of Cayley configuration spaces of 2D unit-distance grid graphs mentioned in Sect. 6.1.6, [11] presents partial results on the configuration spaces of a class of Zeolite sequences.

Acknowledgements Special thanks to the participants of the Barbados (McGill) workshops on Rigidity and molecular modeling organized by Ileana Streinu and Mike Thorpe in 2006 and 2007. Those workshops helped in discovering several of the applications listed in Sect. 6.2.

References

1. Alfakih, A.Y., Khandani, A., Wolkowicz, H.: Solving Euclidean distance matrix completion problems via semidefinite programming. *Comput. Optim. Appl.* **12**, 13–30 (1999)
2. Bădoiu, M., Dhamdhere, K., Gupta, A., Rabinovich, Y., Räcke, H., Ravi, R., Sidiropoulos, A.: Approximation algorithms for low-distortion embeddings into low-dimensional spaces. In: *SODA '05: Proceedings of the 16th annual ACM-SIAM symposium on Discrete algorithms*, pp. 119–128 (2005)
3. Belk (Sloughter), M., Connelly, R.: Realizability of graphs. *Discrete Comput. Geom.* **37**(2), 125–137 (2007)
4. Belk, M.: Realizability of graphs in three dimensions. *Discrete Comput. Geom.* **37**(2), 139–162 (2007)
5. Biswas, P., Lian, T.-C., Wang, T.-C., Ye, Y.: Semidefinite programming based algorithms for sensor network localization. *ACM Trans. Sen. Netw.* **2**(2), 188–220 (2006)
6. Blumenthal, L.M.: *Theory and Applications of Distance Geometry*. Oxford University, London (1953)
7. Borcea, C., Sitharam, M., Streinu, I.: The admissible distance polytope of partial 2-trees (in preparation)
8. Cayley, A.: A theorem in the geometry of position. *Camb. Math. J.* **II**, 267–271 (1841)
9. Connelly, R.: Generic global rigidity. *Discrete Comput. Geom.* **33**(4), 549–563 (2005)
10. Crippen, G.M., Havel, T.F.: *Distance Geometry and Molecular Conformation*. Chemometrics Series, vol. 15. Research Studies Press, Somerset (1998)
11. Gao, H.: *Geometric under-constraints*. Ph.D. Thesis, University of Florida, Gainesville, FL (2008)
12. Gao, H., Sitharam, M.: Combinatorial classification of 2D underconstrained systems. In: Sung-il Pae and Hyungju Park (eds.), *Proceedings of the Seventh Asian Symposium on Computer Mathematics (ASCM 2005)*, pp. 118–127 (2005)
13. Gao, H., Sitharam, M.: Characterizing 1-Dof Henneberg-I graphs with efficient configuration spaces. [arXiv:0810.1997](https://arxiv.org/abs/0810.1997) [cs.CG]
14. Hausman, J.-C.: Geometric descriptions of polygon and chain spaces. *Topology and Robotics. Am. Math. Soc.* **438**, 47–57 (2007)
15. Hoffmann, C.M., Lomonosov, A., Sitharam, M.: Decomposition of geometric constraints systems, part i: performance measures. *J. Comput.* **31**(4), 367–408 (2001)
16. Hoffmann, C.M., Lomonosov, A., Sitharam, M.: Decomposition of geometric constraints systems, part ii: new algorithms. *J. Symb. Comput.* **31**(4), 409–427 (2001)
17. Gortler, S.J., Healy, A.D., Thurston, D.P.: Characterizing generic global rigidity. [arXiv:0710.0926v3](https://arxiv.org/abs/0710.0926v3) [math.MG]
18. Gower, T.: Weblog. <http://gowers.wordpress.com/2007/09/11/what-might-an-expository-mathematical-wiki-be-like/>
19. Graver, J.E., Servatius, B., Servatius, H.: *Combinatorial Rigidity*. Graduate Studies in Math. AMS, Providence (1993)
20. Hendrickson, B.: Conditions for unique graph realizations. *SIAM J. Comput.* **21**(1), 65–84 (1992)
21. Jackson, B., Jordán, T.: Connected rigidity matroids and unique realizations of graphs. *J. Comb. Theory, Ser. B* **94**(1), 1–29 (2005)
22. Jackson, B., Jordán, T., Szabadka, Z.: Globally linked pairs of vertices in equivalent realizations of graphs. *Discrete Comput. Geom.* **35**(3), 493–512 (2006)
23. Jacobs, D.J., Hendrickson, B.: An algorithm for two dimensional rigidity percolation: The pebble game. *J. Comput. Phys.* **137**, 346–365 (1997)

24. Joan-Arinyo, R., Soto-Riera, A., Vila-Marta, S., Vilaplana-Pasto, J.: Transforming an under-constrained geometric constraint problem into a well-constrained one. In: *Symposium on Solid Modeling and Applications 2003*, pp. 33–44 (2003)
25. Kempe, A.B.: On a general method of describing plane curves of the n th degree by linkwork. *Proc. Lond. Math Soc.* **7**, 213–216 (1876)
26. Laman, G.: On graphs and rigidity of plane skeletal structures. *J. Eng. Math.* **4**, 331–340 (1970)
27. Lee, A., Streinu, I., Theran, L.: Finding and maintaining rigid components. In: *17th Canadian Conference on Computational Geometry*, pp. 219–222 (2005)
28. van der Meiden, H.A., Bronsvort, W.F.: A constructive approach to calculate parameter ranges for systems of geometric constraints. *Comput. Aided Des.* **38**(4), 275–283 (2006)
29. Menger, K.: New foundation for Euclidean geometry. *Am. J. Math.* **53**, 721–745 (1931)
30. Owen, J.C., Power, S.C.: Algebraic solution for geometry from dimensional constraints. In: *ACM Symp. Found. of Solid Modeling*, Austin, TX, pp. 397–407 (1991)
31. Owen, J.C., Power, S.C.: The nonsolvability by radicals of generic 3-connected planar graphs. *Trans. AMS* **359**(5), 2269–2303 (2006)
32. Saxe, J.B.: Embeddability of weighted graphs in k -space is strongly NP-hard. In: *Proc. 17th Allerton Conf. in Communications, Control, and Computing*, pp. 480–489 (1979)
33. Schoenberg, I.J.: Metric spaces and positive definite functions. *Trans. Am. Math. Soc.* **44**(3), 522–536 (1938)
34. Sitharam, M.: Graph based geometric constraint solving: problems, progress and directions. In: D. Dutta, R. Janardhan, and M. Smid (eds.), *AMS-DIMACS Volume on Computer Aided Design* (2005)
35. Sitharam, M., Peters, J., Gao, H., Kurnikova, M.: Exact and Efficient description of helix packing configuration space (in preparation)
36. man-Cho So, A., Ye, Y.: A semidefinite programming approach to tensegrity theory and realizability of graphs. In: *SODA*, pp. 766–775 (2006)
37. Zhang, G.F., Gao, X.S.: Well-constrained completion and decomposition for under-constrained geometric constraint problems. *Int. J. Comput. Geom. Appl.* **16**(5–6), 461–478 (2006)