THE GEOMETRY AND TOPOLOGY OF RECONFIGURATION

R. GHRIST AND V. PETERSON

ABSTRACT. A number of reconfiguration problems in robotics, biology, computer science, combinatorics, and group theory coordinate local rules to effect global changes in system states. We define for any such reconfigurable system a cubical complex — the **state complex** — which coordinates independent local moves. We prove classification and realization theorems for state complexes, using CAT(0) geometry as the primary tool. We also classify the topology of spaces of optimal reconfiguration paths using techniques from CAT(0) geometry.

1. INTRODUCTION

There are many contexts in which one wants to control in an optimal fashion systems which are best described as **reconfigurable**. Manufacturing problems possessing large numbers of non-sequential assembly steps provide one class of examples, as do certain problems in theoretical computer science: e.g., asynchronous processors with shared read/write memory. Various approaches for representing such systems include Petri nets [35], high-dimensional automata [31], and process graphs [27]. We initiate a more geometric/topological approach, influenced by ideas from geometric group theory and Alexandrov geometry.

This perspective is entirely natural. Indeed, the initial step in a representation of a complicated reconfigurable system is to represent it as a combinatorial **transition graph** whose vertices represent states and whose edges represent elementary transitions from one state to the next (as in the robotics example of [10]). This is analogous to the construction of a Cayley graph for a group presentation, with the primary difference being that a transition graph may not be homogeneous — not all 'generators' may be applicable at any given state.

We extend the notion of a transition graph in a natural manner, by regarding the graph as the 1-dimensional skeleton of a higher dimensional cell complex. In particular, we use moves which are physically independent (or, more suggestively, 'commutative') to define higher dimensional cubes. The result is the **state complex**: a cubical complex which coordinates independent moves in a reconfigurable system. The idea of using cubes to represent concurrent operations goes back at least

1991 *Mathematics Subject Classification*. Primary: 68Q85 ; Secondary: 57Q05.

Key words and phrases. CAT(0) geometry, configuration space, nonpositive curvature.

RG and VP supported by DARPA # HR0011-05-1-0008 and by NSF PECASE Grant # DMS - 0337713.

to Pratt's paper on high-dimensional automata in 1990 [31]. This paper initiated several lines of research into *geometric concurrency* [20, 27, 32, 33]. Our work differs from this in two principal ways. (1) In high-dimensional automata, the edges of the diagram are oriented. As such, the higher dimensional cubes must be given a partial order, and all questions about the topology of these spaces specialize to delicate notions of directed homotopy of directed paths, etc. (2) The tools used for high-dimensional automata are category-theoretic in nature, with the focus being the determination of the correct setting in which to derive topological invariants for directed homotopy equivalence.

Our departure from this work is to apply tools from CAT(0) geometry that yield more global information. The term CAT(0) was coined by Gromov, and expresses the historical reliance on the work of Cartan, Alexandrov, and Toponogov: briefly, a CAT(0) space is one which is completely devoid of positive curvature, as measured by geodesic triangles. The notion of curvature bounds expressed through geodesic triangles is classical. More recently, the impact of CAT(0) geometry on mathematics has been both extensive and deep, especially in the field of geometric group theory (see [6] and references therein for an overview). The particular case of CAT(0) cube complexes has of late received much attention from the geometric group theory community: see e.g., [6, 11, 26, 28, 30, 34].

It is increasingly clear that CAT(0) geometry is of great importance in applications. A principal example of this appears in the paper of Billera, Holmes, and Vogtmann on spaces of phylogenetic trees [5], in which CAT(0) geometry is used to solve problems of qualitative classification in biological systems. Other examples include the precursor to this work [21, 3] on configuration spaces for metamorphic robots, and more recent work [24, 25] on Pareto optimization in robotics. Extending classical results on planar pursuit-evasion to higher dimensional domains is impossible in general but both possible and fairly simple on CAT(0) domains [4].

Section 2 gives definitions of reconfigurable systems and state complexes. Section 3 is a collection of interesting examples of reconfigurable systems and their state complexes, both physical and abstract. These examples lead one to observe the prevalence of [discrete] negative curvature, an intuition that is confirmed in the proof of the local CAT(0) geometry of state complexes in Section 4. The main results of classification and realization are presented there and in Section 5, with observations and unresolved questions given in 6.

The most significant classification results for state complexes in this paper are as follows:

(1) Any reconfigurable system yields a state complex which is locally CAT(0) .
(2) One can realize cubical complexes with arbitrary (flag) link structures as state complexes for reconfigurable systems.
(3) Any locally CAT(0) subcomplex of a product of graphs can be realized as a state complex of a reconfigurable system.

(4) Fundamental groups of compact state complexes embed into Artin right-angled groups and thus are linear.

2. DEFINITIONS

2.1. **Reconfigurable systems.** A reconfigurable system is a collection of states on a graph, where each state is thought of as a vertex labeling function. Any state can be modified by local rearrangements, these local changes being rigidly specified. We distinguish between the amount of information needed to determine the legality of an elementary move (the "support" of the move) and the precise subset on which the reconfiguration physically occurs (the "trace" of the move).

Definition 2.1. Fix $\mathcal{A}$ to be a set of labels. Fix $\mathcal{G}$ to be a graph. A **generator** ϕ for a local reconfigurable system is a collection of three objects:

(1) the **support**, $\text{SUP}(\phi) \subset \mathcal{G}$, a subgraph of $\mathcal{G}$;
(2) the **trace**, $\text{TR}(\phi) \subset \text{SUP}(\phi)$, a subgraph of $\text{SUP}(\phi)$;
(3) an *unordered* pair of **local states**

$$\mathbf{u}_0^{loc}, \mathbf{u}_1^{loc} : V(\text{SUP}(\phi)) \to \mathcal{A},$$

which are labelings of the vertex set of $\text{SUP}(\phi)$ by elements of $\mathcal{A}$. These local states must agree on $\text{SUP}(\phi) - \text{TR}(\phi)$: i.e.,

$$\mathbf{u}_0^{loc}\Big|_{\text{SUP}(\phi)-\text{TR}(\phi)} = \mathbf{u}_1^{loc}\Big|_{\text{SUP}(\phi)-\text{TR}(\phi)}. \tag{2.1}$$

All generators are assumed to be **nontrivial** in the sense that $\mathbf{u}_0^{loc} \neq \mathbf{u}_1^{loc}$.

Definition 2.2. A **state** is a labeling of the vertices of $\mathcal{G}$ by $\mathcal{A}$. A generator ϕ is said to be **admissible** at a state $\mathbf{u}$ if $\mathbf{u}|_{\text{SUP}(\phi)} = \mathbf{u}_0^{loc}$. For such a pair $(\mathbf{u}, \phi)$, we say that the **action** of ϕ on $\mathbf{u}$ is the new state given by

$$\phi[\mathbf{u}] := \begin{cases} \mathbf{u} & : \text{ on } \mathcal{G} - \text{SUP}(\phi) \\ \mathbf{u}_1^{loc} & : \text{ on } \text{SUP}(\phi) \end{cases}, \tag{2.2}$$

Remark 2.3. Since the local states of each generator are unordered, it follows that any generator ϕ which is admissible at a state $\mathbf{u}$ is also admissible at the state $\phi[\mathbf{u}]$, and that $\phi[\phi[\mathbf{u}]] = \mathbf{u}$.

Definition 2.4. A **reconfigurable system** on $\mathcal{G}$ consists of a collection of generators and a collection of states closed under all possible admissible actions.

We identify certain special types of systems.

Definition 2.5. A reconfigurable system is said to be **locally finite** if the number of generators admissible at any state has a finite bound, independent of the state. A reconfigurable system is said to be **homogeneous** if the generators are independent of location, in the following sense. Given any generator ϕ and any graph embedding $\delta : \text{SUP}(\phi) \hookrightarrow \mathcal{G}$, then $\delta \circ \phi$ is also a generator for the system. Local states, and all other features, are then composed with δ.

2.2. **The state complex.**

Definition 2.6. In a reconfigurable system, a collection of generators $\{\phi_{\alpha_i}\}$ is said to **commute** if

$$\text{TR}(\phi_{\alpha_i}) \cap \text{SUP}(\phi_{\alpha_j}) = \emptyset \quad \forall i \neq j. \tag{2.3}$$

Commutativity connotes physical independence. Consider the system of Fig. 1[left], which consists of planar hexagonal cells in a hex lattice. For the moment, think of a 'generator' as representing a hexagon pivoting to an unoccupied neighboring lattice point. It is the case that a pair of commuting generators yields a square in the transition graph of states.

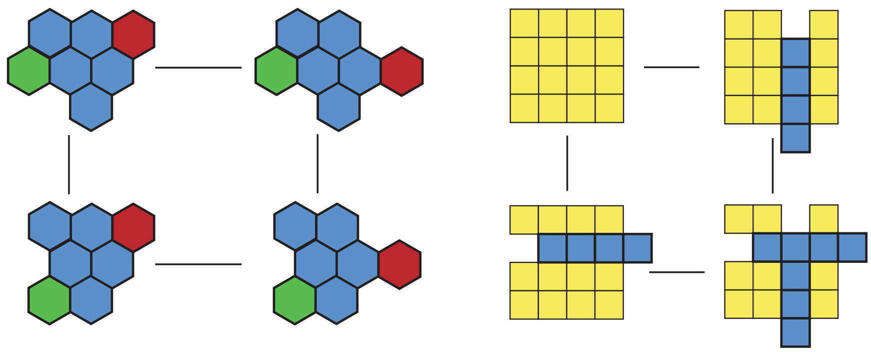

FIGURE 1. Examples of commuting and noncommuting local moves which form a 4-cycle in the transition graph: [left] pivoting hexagons lead to commutative moves; [right] sliding rows/columns which intersect does not commute.

Compare this with a planar sliding block example in Fig. 1 [right]: 'generators' consist of sliding a row or column of squares one unit. Although the pair of moves illustrated forms a square in the transition graph, this particular pair of generators does not commute. Physically, it is obvious why these moves are not independent: sliding the column part-way obstructs sliding a transverse row. If we were to make this a formal reconfigurable system, then we would specify that the trace of a generator is the entire row or column. The traces of the generators illustrated then intersect.

We define the state complex to be the cube complex with an abstract k-cube for each collection of k admissible commuting generators:

Definition 2.7. The **state complex** $\mathcal{S}$ of a local reconfigurable system is the following abstract cubical complex. Each abstract k-cube $e^{(k)}$ of $\mathcal{S}$ is an equivalence class $[\mathbf{u}; (\phi_{\alpha_i})_{i=1}^k]$ where

(1) $(\phi_{\alpha_i})_{i=1}^k$ is a k-tuple of commuting generators;
(2) $\mathbf{u}$ is some state for which all the generators $(\phi_{\alpha_i})_{i=1}^k$ are admissible; and

(3) $[\mathbf{u}_0; (\phi_{\alpha_i})_{i=1}^k] = [\mathbf{u}_1; (\phi_{\beta_i})_{i=1}^k]$ if and only if the list (β_i) is a permutation of (α_i) and $\mathbf{u}_0 = \mathbf{u}_1$ on the set $\mathcal{G} - \bigcup_i \text{SUP}(\phi_{\alpha_i})$.

The boundary of each abstract k-cube is the collection of $2k$ faces obtained by deleting the i^{th} generator from the list and using $\mathbf{u}$ and $\phi_{\alpha_i}[\mathbf{u}]$ as the ambient states, for $i = 1 \ldots k$. Specifically,

$$\partial[\mathbf{u}; (\phi_{\alpha_i})_{i=1}^k] = \bigcup_{i=1}^k \left([\mathbf{u}; (\phi_{\alpha_j})_{j\neq i}] \cup [\phi_{\alpha_i}[\mathbf{u}]; (\phi_{\alpha_j})_{j\neq i}] \right) \tag{2.4}$$

The weak topology is used for reconfigurable systems which are not locally finite. In the locally finite case, the state complex is a locally compact cubical complex.

It follows from repeated application of Remark 2.3 that the k-cells are well-defined with respect to admissibility of actions. The following two lemmas are trivial and included only as an exercise in using the definitions.

Lemma 2.8. *(a) The 0-skeleton of $\mathcal{S}$, $\mathcal{S}^{(0)}$, is the set of states in the reconfigurable system. (b) The 1-skeleton of $\mathcal{S}$, $\mathcal{S}^{(1)}$, is precisely the transition graph.*

PROOF: (a) Vertices of $\mathcal{S}$ consist of equivalence classes consisting of zero (i.e., no) actions of generators up to permutation, together with a state defined on the complement of the supports of the actions. As there are no actions, each 0-cell is precisely a single state of the reconfigurable system.

(b) A 1-cell of $\mathcal{S}$ is an equivalence class of the form $[\mathbf{u}; (\phi)]$. The only other representative of the equivalence class is $[\phi[\mathbf{u}]; (\phi)]$; hence, the 1-cells are precisely the edges in the transition graph. Clearly, the boundary of $[\mathbf{u}; (\phi)]$ is the pair of 0-cells $[\mathbf{u}; (\cdot)]$ and $[\phi[\mathbf{u}]; (\cdot)]$. □

Lemma 2.9. *In any state complex $\mathcal{S}$, the closure of each k-cell in $\mathcal{S}$ is an embedded k-cube.*

PROOF: Assume to the contrary that two vertices $\mathbf{u}_0$ and $\mathbf{u}_1$ of a k-cube are equal. By choosing the cube of smallest dimension containing $\mathbf{u}_0$ and $\mathbf{u}_1$, it may be assumed that these are antipodal vertices in the k-cube. Thus,

$$\mathbf{u}_0 = \mathbf{u}_1 = \phi_{\alpha_k}\left[\phi_{\alpha_{k-1}}[\cdots[\phi_{\alpha_1}[\mathbf{u}_0]\cdots]\right].$$

By Definition 2.6, $\text{TR}(\phi_{\alpha_i})$ are all disjoint. Hence,

$$\mathbf{u}_1|_{\text{TR}(\phi_{\alpha_i})} = \phi_{\alpha_i}[\mathbf{u}_0]|_{\text{TR}(\phi_{\alpha_i})},$$

contradicting the fact that generators are nontrivial. □

Based on Lemma 2.8, we say that two reconfigurable systems are **isomorphic** if their state complexes are isomorphic (there is a homeomorphism between them preserving the cubical structure).

Lemma 2.10. *Any reconfigurable system is isomorphic to a homogeneous reconfigurable system.*

PROOF: Starting with an inhomogeneous reconfigurable system with domain $\mathcal{G}$ having vertex set $V(\mathcal{G})$ and alphabet $\mathcal{A}$, change the alphabet to be $\mathcal{A} \times V(\mathcal{G})$. Modify the labels of the corresponding generators to incorporate vertex labels, increasing the number of generators if necessary. This can be done so that generators are applicable only at the specified vertices. □

3. EXAMPLES

In many of the examples which follow, the 'states' correspond to some collection of geometric cells, edges, or other physical objects: e.g., Fig. 1. We will occasionally be imprecise in specifying what the underlying graph $\mathcal{G}$ for the system is — in such cases, one should use the dual graph to the collection of cells/edges.

Example 3.1 (hex-lattice metamorphic robots). This reconfigurable system is based on the first metamorphic robot system pioneered by Chirikjian [10]. It consists of a finite aggregate of planar hexagonal units locked in a hex lattice, with the ability to pivot sufficiently unobstructed units on the boundary of the aggregate.

More specifically, $\mathcal{G}$ is a graph whose vertices correspond to hex lattice points and whose edges correspond to neighboring lattice points. The alphabet is $\mathcal{A} = \{0, 1\}$ with 0 connoting an unoccupied site and 1 connoting an occupied site. There is one type of generator, represented in Fig. 2[left], which generates a homogeneous system: this local rule can be applied to any translated or rotated position in the lattice. This generator allows for local changes in the topology of the aggregate (disconnections are possible). For physical systems in which this is undesirable — say, for power transmission purposes — one can choose a generator with larger support.

As an example of a state complex for this system, consider a workspace $\mathcal{G}$ consisting of three rows of lattice points with a line of occupied cells as in Fig. 3. This line of cells can "climb" on itself from the left and migrate to the right, one by one. The entire state complex is illustrated in Fig. 3[center]. Although the transition graph appears complicated, this state complex is contractible and remains so for any length channel.

Example 3.2 (2-d articulated planar robot arm). Consider as a domain $\mathcal{G}$ the lattice of edges in the first quadrant of the plane. This system consists of two types of generators, pictured in Fig. 4. The support of each generator is the union of eight edges as shown. The trace of each generator is as described in the figure caption. Beginning with a state having N vertical edges end-to-end, the reconfigurable system models the position of an articulated robotic arm which is fixed at the origin and which can (1) rotate at the end and (2) flip corners as per the diagram. This arm is **positive** in the sense that it may extend up and to the right only.

The state complex in the case $N = 5$ is illustrated in Fig. 5. Note that there can be at most three independent motions (when the arm is in a "staircase" configuration); hence the state complex has top dimension three. In this case also, although

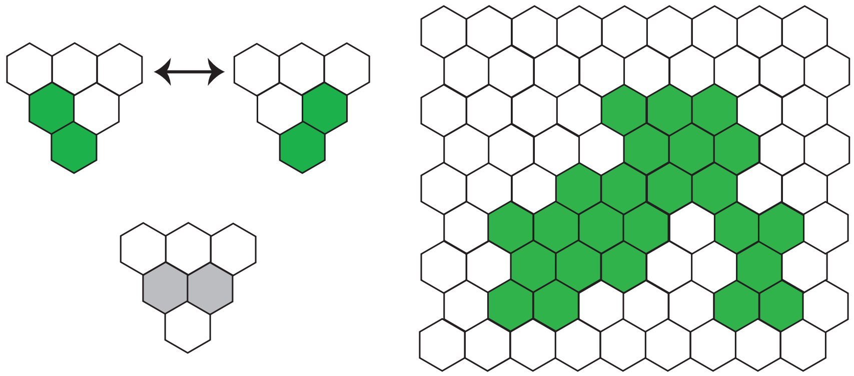

FIGURE 2. The generator for a 2-d hexagonal lattice system with pivoting locomotion. The domain is the graph dual to the hex lattice shown. Shaded cells are occupied, white are unoccupied. [left, top] The local states $\mathbf{u}_0^{loc}$ and $\mathbf{u}_1^{loc}$ are shown. [left, bottom] The support of the generator, with trace shaded. [right] A typical state in this reconfigurable system.

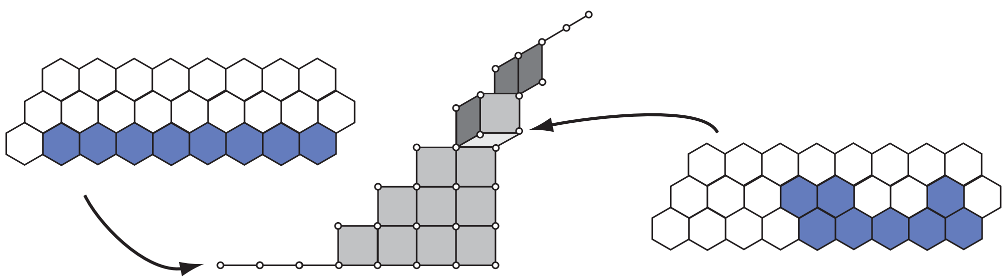

FIGURE 3. For a line of hexagons filing out of a constrained tunnel, the state complex is contractible.

the transition graph for this system is complicated, the state complex itself is contractible: this is the case for all lengths N.

Example 3.3 (configuration space of points on a graph). Consider a graph $\mathcal{G}$ and alphabet $\mathcal{A} = \{0, \ldots, n\}$ used to specify empty/occupied vertices. There are n types of generators $\{\phi_i\}_1^n$ in this homogeneous system, one for each nonzero element of $\mathcal{A}$. The support and trace of each ϕ_i is precisely the closure of an (arbitrary) edge. The local states of this ϕ_i evaluate to 0 on one of the endpoints and i on the other. The homogeneous reconfigurable system generated from a state $\mathbf{u}$ on $\mathcal{G}$ having exactly one vertex labeled i for each $i = 1, \ldots, n$ mimics an ensemble of N distinct non-colliding points on the graph $\mathcal{G}$. If we reduce the alphabet to $\{0, 1\}$, then the system represents n identical agents.

This system is a discrete model of a collection of robots which are constrained to travel along tracks or guidewires [22, 23]. The associated state complexes for these

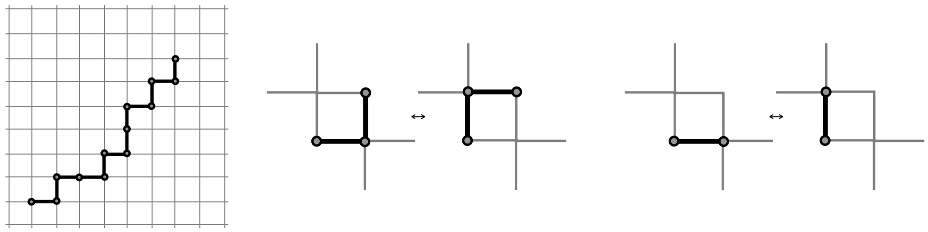

FIGURE 4. A positive articulated robot arm example [left] with fixed endpoint. One generator [center] flips corners and has as its trace the central four edges. The other generator [right] rotates the end of the arm, and has trace equal to the two activated edges.

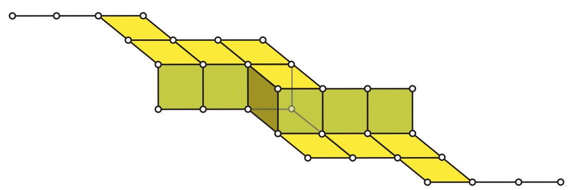

FIGURE 5. The state complex of a 5-link positive arm has one cell of dimension three, along with several cells of lower dimension.

systems is a discrete type of configuration space for these systems. Such spaces were considered independently by Abrams [1] and also by Swiatkowski [38].

For example, if the graph is K_5 (the complete graph on five vertices), $N = 2$, and $\mathcal{A} = \{0, 1, 2\}$, it is straightforward to show that each vertex has a neighborhood with six edges incident and six 2-cells patched cyclically about the vertex. Therefore, $\mathcal{S}$ is a closed surface. One can (as in [2]) count that there are 20 vertices, 60 edges, and 30 faces in the state complex. The Euler characteristic of this surface is therefore -10. This surface can be given an orientation; thus, the state complex has genus six.

Example 3.4 (digital microfluidics). An even better physical instantiation of the previous system arises in digital microfluidics [17, 18]. In this setting, small (e.g., 1mm diameter) droplets of fluid can be quickly and accurately manipulated on a plate covering a network of current-controlled wires by an electrowetting process that exploits surface tension effects to propel a droplet. Applying a current drives the droplet a discrete distance along the wire. In this setting, one desires a "laboratory on a chip" in which droplets of various chemicals can be positioned, mixed, and then directed to the appropriate outputs.

Representing system states as marked vertices on a graph is appropriate given the discrete nature of the motion by electrophoresis on a graph of wires. This adds a

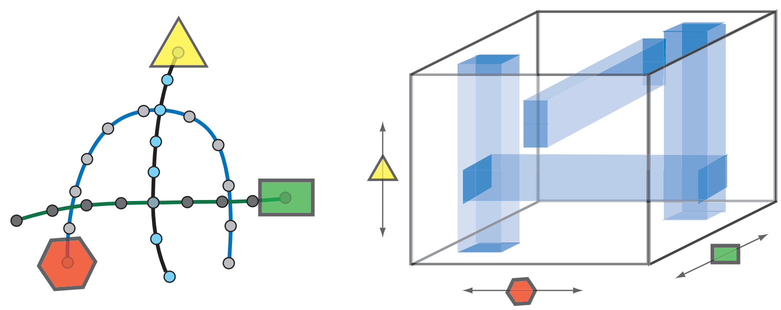

FIGURE 6. [left] A coordination problem with three robots translating on (discretized, intersecting) intervals. [right] The state complex approximation to the coordination space, with collision set shaded.

few new ingredients to the setting of the previous example, though. For n different chemical agents, an alphabet of $\{0, \ldots, n\}$ is appropriate (the '0' connoting absence); however, a typical state may have many vertices with the same nonzero label (corresponding to the number of droplets of substance i in use at a given time). Furthermore, it is possible to mix droplets by merging them together, rapidly oscillating along an edge, then splitting the mixed product. This leads to a new type of generator of the form $(i - j) \Longleftrightarrow (k - k)$.

Example 3.5 (robot coordination). There is a broad generalization of configuration spaces of graphs developed in [25, 24] which has an interpretation as a state complex. We outline a simple example. Consider a collection of N planar graphs $(\Gamma_i)_1^N$, each embedded in the plane of a common workspace (with intersections between different graphs permitted). On each Γ_i, a robot R_i with some particular fixed size/shape is free to translate along Γ_i: one thinks of the graph as being a physical groove in the floor, or perhaps an electrified overhead guidewire. The **coordination space** of this system is defined to be the space of all configurations in $\prod_i \Gamma_i$ for which there are no collisions — the robots R_i have no intersections.

We can approximate these coordination spaces by the following reconfigurable systems. Assume that each graph Γ_i has been refined by adding multiple (trivial) vertices along the interiors of edges. We will approximate the robot motion by performing discrete jumps to neighboring vertices, much as in Example 3.3.

Let the underlying graph be $\mathcal{G} := \coprod \Gamma_i$, the disjoint union of the individual graphs. The generators for this system are as follows. For each edge $\alpha \in E(\Gamma_i)$, there is exactly one generator ϕ_α. The trace is the edge itself, $\text{TR}(\phi_\alpha) = \alpha$, and the generator corresponds to sliding the robot R_i from one end of the edge to the other. The support, $\text{SUP}(\phi_\alpha)$ consists of the edge $\alpha \in E(\Gamma_i)$ along with any other edges β in Γ_j $(j \neq i)$ for which the robot R_j sliding along the edge β can collide with R_i as it slides along α. The alphabet is $\mathcal{A} = \{0, 1\}$ and the local states for ϕ_α have zeros at all vertices of all edges in the support, except for a single 1 at the boundary vertices of α (these two boundary vertices yield the two local states, as in Example 3.3).

Any state for this reconfigurable system is one for which all vertices of each Γ_i are labeled with zeros except for one vertex with a label 1. The resulting state complex is a cubical complex which approximates the cylindrical coordination space, as pictured in Fig. 6 [25, 24]. Of course, in the case where $\Gamma_i = \Gamma$ for all i and the robots R_i are sufficiently small, this reconfigurable system is exactly that of Example 3.3.

Example 3.6 (protein folding). Certain discrete models of protein folding are amenable to a reconfigurable system analysis. In particular, the model proposed by Sali et al. [36, 37] treats the protein molecule as a piecewise-linear chain in a cubic lattice of edges. They model the folding process as a sequence of applications of local rules (see Fig. 1(b) of [37]) reminiscent of the articulated robot arm of Example 3.2.

It is especially simple to write a reconfigurable system for a closed-chain version of the model in [36, 37]. One represents the protein chains as states in a cubic edge lattice with alphabet $\{0, 1\}$ (occupied vs. unoccupied edges). Generators correspond to the local rules of Fig. 7 which flip segments of length two and three respectively. The resulting state complex will be a cubical complex approximating the configuration space of the model protein loop.

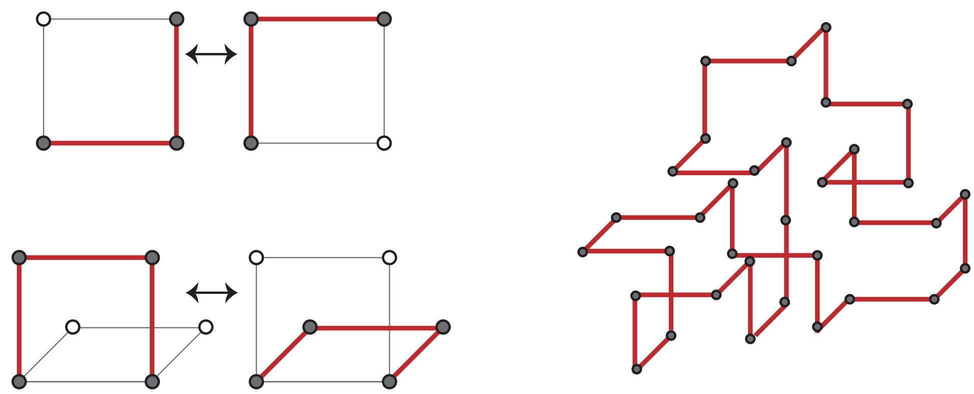

FIGURE 7. Two local moves [left] for a simple model of a closed-chain protein [right] rotate either two or three consecutive edges to change conformations.

Example 3.7 (permutohedra). Consider a graph $\mathcal{G}$ and any finite alphabet. The generators have support and trace equal to a single edge; the local states exchange distinct labels on the two vertices of the edge. This example becomes the system of Example 3.3 if one permits only exchanges between the label 0 (*i.e.*, unoccupied sites) and any of the other non-zero labels (*i.e.*, occupied sites).

The geometry of the state complex is very clean in cases where the graph $\mathcal{G}$ is a 5-gon and the alphabet consists of five elements, one per vertex. Since $\mathcal{G}$ has five vertices, there are $5! = 120$ vertices in the state complex, and these states may uniquely be identified with S_5, the permutation group on 5 elements. The generators of the reconfigurable system are adjacent transpositions, which comprise the

Coxeter presentation of S_5. Each vertex has a neighborhood with five generators (one for each edge of $\mathcal{G}$). Since each edge of $\mathcal{G}$ is disjoint from exactly two other edges of $\mathcal{G}$ (which are not, themselves, disjoint), the neighborhood of the vertex is a cyclic arrangement of five squares, as in Fig. 8. The state complex is thus a closed 2-manifold. One can explicitly specify a global orientation for this 2-complex. Given 120 vertices with this local geometry implies an Euler characteristic of -30. The state complex is thus a closed orientable surface of genus 16, which 'fills in' the Cayley graph for the Coxeter presentation of S_5.

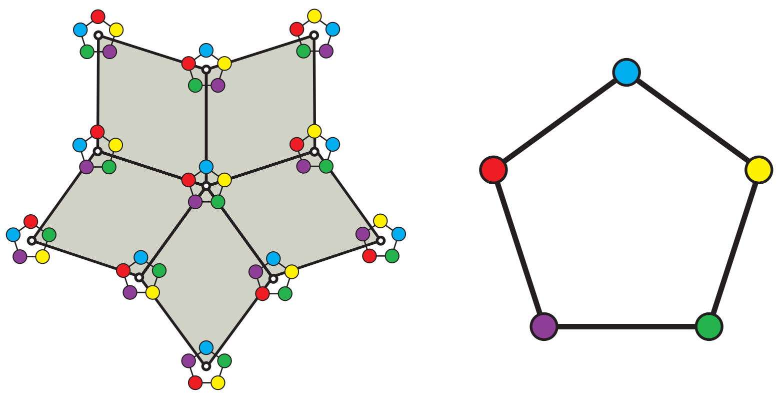

FIGURE 8. Generators which permute neighboring vertices on a 5-gon [right] leads to a state complex each vertex of which has a 5-gon link [left]. A small picture of the state in superimposed on vertices of the state complex.

In the case of $\mathcal{G} = K_5$, the state complex is two-dimensional, but is not a manifold. Each vertex in $\mathcal{S}$ is surrounded by ten edges and fifteen 2-cells. We leave it to the reader to check that the link of each vertex in $\mathcal{S}$ is the Petersen graph.

For reasons of space, we omit many of the other interesting examples of reconfigurable systems and their state complexes. Some relating to robotics applications can be found in [3]. The work of Farley [19] gives a cubical complex structure for the diagram group of a semigroup presentation, which may be realized as a state complex. Spaces of triangulations of polygonal domains with edge-flips as generators also form an interesting example related to associahedra. Finally, the spaces of rooted labeled trees used by Billera, Holmes, and Vogtmann [5] to work with phylogenetic trees are realizable as state complexes for a reconfigurable system (see Corollary 5.14). We suspect there are many additional interesting examples.

4. Geometry

From the previous section, one observes the prevalence of state complexes which are either contractible or have contractible universal cover (as is the case for closed surfaces of nonzero genus). The answer to the question of whether this holds in general depends on geometric properties of state complexes.

Remark 4.1. State complexes inherit a natural piecewise Euclidean geometry from the generators. Each generator ϕ_i corresponds to some "move" which, we assume, can be executed at some uniform speed and requires L_i time units to do so. This gives a natural linear metric to the edges of the transition graph which correspond to ϕ_i: such edges have length L_i. Since higher dimensional cells of $\mathcal{S}$ are determined by concurrent executions, these cubes inherit a natural flat product metric. The result is that k-cells of $\mathcal{S}$ are Euclidean rectangular prisms. We note that in many examples, L_i is independent of i, and the resulting metric on $\mathcal{S}$ has all cells Euclidean unit cubes. We will call $\mathcal{S}$ a piecewise-Euclidean cubical complex, even in cases where the edge lengths vary. For the remainder of this paper, we work under the natural assumption that the set $\{L_i\}$ of lengths is bounded away from both zero and infinity.

4.1. **Curvature for cubical complexes.**

Piecewise Euclidean cubical complexes are flat in the interiors of the cubical cells; however, non-zero curvature can be concentrated at places where several cells meet. For example, a surface built from flat 2-cells can be seen to have a discrete curvature which depends on the number of 2-cells incident to a vertex. The case of four incident cells implies zero curvature; that of three cells implies positive curvature; and that of five or more cells implies negative curvature. A broad extension of curvature to general metric spaces is made precise in the classical work of Alexandrov and others, in which triangles with geodesic edges are used to measure curvature bounds. A geodesic path in X is a rectifiable path whose length is equal to the metric distance between the points.

Assume that X is a metric space for which geodesic paths exist (the piecewise Euclidean cubical complexes we work with here all have this property [6]). Consider any triangle T in X with geodesic edges of length a, b, and c. Build a **comparison triangle** T' in the Euclidean plane whose sides also have length a, b, and c respectively. Choose a geodesic chord of T and measure its length d. In T', measure the length d' of the associated chord.

Definition 4.2. A metric space X is **CAT(0)** if for every geodesic triangle T it holds that $d \leq d'$ for all chords of T. One says that X is **nonpositively curved** (or **NPC**) if X is locally CAT(0) ; that is, if $d \leq d'$ for all sufficiently small T.

A space is CAT(0) if and only if it is simply connected and NPC. Being NPC implies a variety of topological consequences: for example, the universal cover is contractible and the fundamental group is torsion-free. See, *e.g.*, [6] for a thorough introduction to spaces of nonpositive curvature.

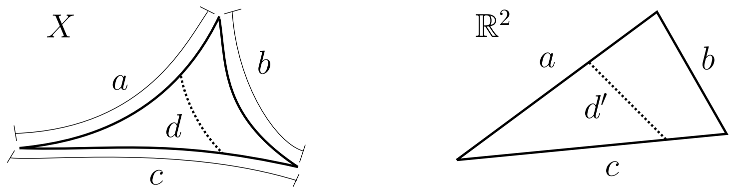

FIGURE 9. Comparison triangles measure curvature bounds.

4.2. **The link condition.** There is a well-known combinatorial approach to determining when a cubical complex is nonpositively curved due to Gromov.

Definition 4.3. Let X denote a cell complex and let v denote a vertex of X. The **link** of v, $\ell k[v]$, is defined to be the abstract simplicial complex whose k-dimensional simplices are the $(k+1)$-dimensional cells incident to v with the natural boundary relationships.

Certain global topological features of a metric cubical complex are completely determined by the local structure of the vertex links: a theorem of Gromov [26] asserts that a finite dimensional Euclidean cubical complex is NPC if and only if the link of every vertex is a flag complex without digons. Recall: a **digon** is a pair of vertices connected by two edges, and a **flag complex** is a simplicial complex which is maximal among all simplicial complexes with the same 1-dimensional skeleton. Gromov's theorem permits us an elementary proof of the following general result.

Theorem 4.4. *The state complex of any locally finite reconfigurable system is NPC.*

PROOF: Gromov's theorem is stated for finite dimensional Euclidean cubical complexes with unit length cubes. It holds, however, for non-unit length cubes when there are a finite number of isometry classes of cubes (the *finite shapes* condition) [6]. Locally finite reconfigurable systems possess locally finite and finite dimensional state complexes, which automatically satisfy the finite shapes condition (locally).

Let $\mathbf{u}$ denote a vertex of $\mathcal{S}$. Consider the link $\ell k[\mathbf{u}]$. The 0-cells of the $\ell k[\mathbf{u}]$ correspond to all edges in $\mathcal{S}^{(1)}$ incident to $\mathbf{u}$; that is, actions of generators based at $\mathbf{u}$. A k-cell of $\ell k[\mathbf{u}]$ is thus a commuting set of $k+1$ of these generators based at $\mathbf{u}$.

We argue first that there are no digons in $\ell k[\mathbf{u}]$ for any $\mathbf{u} \in \mathcal{S}$. Assume that ϕ_1 and ϕ_2 are admissible generators for the state $\mathbf{u}$, and that these two generators correspond to the vertices of a digon in $\ell k[\mathbf{u}]$. Each edge of the digon in $\ell k[\mathbf{u}]$ corresponds to a distinct 2-cell in $\mathcal{S}$ having a corner at $\mathbf{u}$ and edges at $\mathbf{u}$ corresponding to ϕ_1 and ϕ_2. By Definition 2.7, each such 2-cell is the equivalence class $[\mathbf{u}; (\phi_1, \phi_2)]$: the two 2-cells are therefore equivalent and not distinct.

To complete the proof, we must show that the link is a flag complex. The interpretation of the flag condition for a state complex is as follows: if at $\mathbf{u} \in \mathcal{S}$, one has a set of k generators ϕ_{α_i}, of which each pair of generators commutes, then the full set of k generators must commute. The proof follows directly from the definitions, especially from two observations from Definition 2.6: (1) commutativity of a set of actions is independent of the states implicated; and (2) any collection of pairwise commutative actions is totally commutative. □

5. REALIZATION

In this section, we work toward a classification of state complexes.

5.1. **Realizing links.** As we have demonstrated in Section 3, it is possible to construct reconfigurable systems whose state complexes are surfaces with negative curvature at each vertex. We extend this class of examples significantly. The following result parallels a well-known theorem of M. Davis [13]; the formalism of reconfigurable systems yields a simple, clean proof.

Theorem 5.1. *Let L be any finite simplicial flag complex. There exists a finite reconfigurable system whose state complex has the property that $\ell k[v] = L$ for all vertices $v \in \mathcal{S}$.*

PROOF: The proof is explicit. Define $\mathcal{G}$ to be the 1-skeleton of L with alphabet $\mathcal{A} = \{0, 1\}$. There is one type of generator per vertex $v \in L$. Its trace is v and its support is equal to v together with the maximal subgraph of $\mathcal{G}$ whose vertices are all more than one edge away from v. The two local states differ only on v. There is one such generator for each possible labeling of SUP − TR. From the definitions, two generators with traces v and v' commute if and only if there is an edge in L between v and v'. Each vertex of $\mathcal{S}$ therefore has link L, since commutativity is determined pairwise and L is flag. □

Example 5.2. One constructs an n-manifold state complex as follows. Consider a simplicial, flag, $(n-1)$-sphere L with C_k simplices of dimension k. If one builds the state complex with link L as in the proof of Theorem 5.1, there are exactly 2^{C_0} vertices in $\mathcal{S}$. A careful count yields that there are exactly $C_{k-1}2^{C_0-k}$ cubes of dimension k in $\mathcal{S}$ (using the convention that $C_{-1} = 1$). Hence, the Euler characteristic of $\mathcal{S}$ is

$$\chi(\mathcal{S}) = \sum_{i=0}^{n} (-1)^i C_{i-1} 2^{C_0-i} \tag{5.1}$$

For example, when $n = 2$, the link is a simplicial circle with $C_0 = C_1 = C \geq 4$ and Euler characteristic $\chi = 2^{C-2}(4-C)$: this number takes on infinitely many different values in C. All of these are negative (reflecting the nonpositive curvature) except for the case $C = 4$: see Fig. 10 for this example.

5.2. **Hyperplanes.** Our proofs rely on notions of **hyperplanes** as developed in [30, 34].

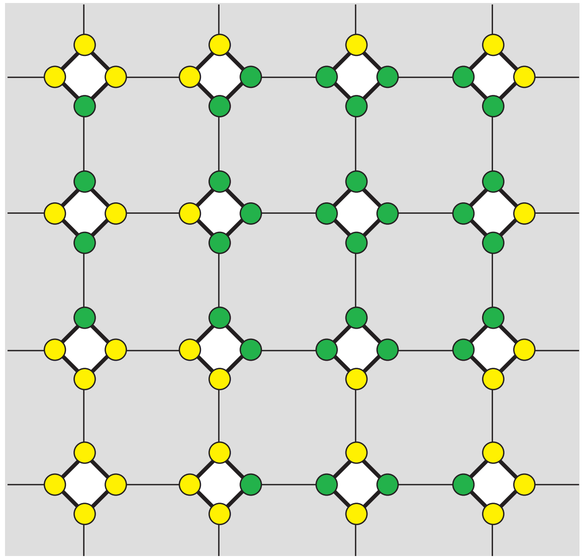

FIGURE 10. A pictorial representation of the proof of Theorem 5.1 in the case where the link is a 4-gon. The state complex $\mathcal{S}$ consists of 16 vertices and 16 squares which together form a 2-torus. Shown is a cut-open version of this torus with each vertex of $\mathcal{S}$ replaced by a copy of the labeled 4-gon link. Edges in $\mathcal{S}$ represent elementary changes in the labels on the link.

Definition 5.3. Let X be a cubical complex, each cube outfitted with coordinates $\{x_i \in [-1, 1]\}$. A **midplane** of a cube $[-1, 1]^k$ is a codimension-1 coordinate plane of the form $\{x_i = 0\}$. Two midplanes M and N in a cubical complex X are said to be **hyperplane equivalent** if there is a sequence of midplanes $M = M_1, M_2, \ldots, M_n = N$ in X such that $M_i \cap M_{i+1}$ is a midplane for every $i = 1, 2, \ldots, (n-1)$. With respect to this equivalence, a **hyperplane** is an equivalence class of midplanes.

We will usually consider a hyperplane as the union of the midplanes in its equivalence class, the midplanes having been glued together via restrictions of the gluing maps for cubes in the complex. The following simple lemma asserts that hyperplanes are dual to generators.

Lemma 5.4. *Let $\mathcal{H}$ denote a hyperplane of a state complex $\mathcal{S}$. Every edge of $\mathcal{S}$ intersecting $\mathcal{H}$ corresponds to the action of a fixed generator $\phi_{\mathcal{H}}$.*

PROOF: The result holds for each midplane of a k-dimensional cube in $\mathcal{S}$. Since the gluing maps for midplane equivalence are the restrictions of the gluing maps for the cubical complex $\mathcal{S}$, this unique generator is transported to each midplane in the hyperplane. □

This generalizes to the following result, stated in terms of carriers. Recall that the **carrier** of a subset U of a cell complex X is $C(U)$, the smallest closed subcomplex of X containing U.

Lemma 5.5. *In any state complex $\mathcal{S}$, the carrier of any hyperplane $\mathcal{H}$ is a cube complex isomorphic to $\mathcal{H} \times [-1, 1]$ with $\mathcal{H}$ corresponding to the zero-section.*

PROOF: The carrier of the hyperplane $C(\mathcal{H})$ is equal to the union of the carriers of the midplanes, and each midplane is equal to the zero-section of its carrier cube.

Assume first that two disjoint but equivalent midplanes in a given hyperplane have carriers which intersect. Since the carrier is a cube complex, there must be two distinct edges in $C(\mathcal{H})$ which are transverse to $\mathcal{H}$ but intersect in a single vertex. Lemma 5.4 implies that these edges correspond to the same generator. Since the edges intersect at a vertex, we have a single generator applied to a state yielding two distinct states: contradiction.

We conclude that $C(\mathcal{H})$ is a bundle over $\mathcal{H}$ with fiber $[-1, 1]$ and zero-section $\mathcal{H}$. If this bundle is nontrivial, then there exists a loop in $\mathcal{H}$ over which the bundle is a Möbius strip. One may assume without loss of generality that this strip lies in the 2-skeleton of $\mathcal{S}$. Choose an arbitrary vertex $\mathbf{u}$ in the strip. Let ϕ_0 denote the generator dual to $\mathcal{H}$ from Lemma 5.4. By nonorientability of the 2-complex, we can express ϕ_0 as the composition

$$\phi_0[\mathbf{u}] = \phi_n \phi_{n-1} \cdots \phi_2 \phi_1 [\mathbf{u}], \tag{5.2}$$

where $\{\phi_i\}_1^n$ are the sequence of (not necessarily distinct) generators which wrap around the strip, each commuting with ϕ_0. By commutativity and (5.2), we have

$$\mathbf{u}|_{\mathrm{TR}(\phi_0)} = \phi_n \cdots \phi_2 \phi_1 [\mathbf{u}]|_{\mathrm{TR}(\phi_0)} = \phi_0[\mathbf{u}]|_{\mathrm{TR}(\phi_0)}. \tag{5.3}$$

This contradicts the fact that ϕ_0 is a nontrivial generator and thus changes its state somewhere on the trace. □

Corollary 5.6. *Not all NPC cubical complexes are realized as the state complex of a reconfigurable system.*

5.3. **Fundamental groups of state complexes.** It is not immediately clear which fundamental groups of NPC cube complexes can arise as the fundamental group of a state complex. We show that fundamental groups of state complexes have some very particular algebraic properties. The following theorem is a mild modification of a proof of Crisp and Wiest [11], or, as well, it follows from a more recent result of Haglund and Wise [28]. Recall that an **Artin right-angled group** is a group with presentation having all relations commutators in the generators. The following result is very satisfying, as it validates the terminology of *generators* and *commuting* from Definitions 2.1 and 2.6.

Theorem 5.7. *The fundamental group of any finite state complex $\mathcal{S}$ embeds into the finitely generated Artin right-angled group whose generators correspond to generators of* *the reconfigurable system and whose commutators correspond to pairs of generators which commute at some state in $\mathcal{S}$.*

The proof follows almost directly from the proof of Theorem 2 of [11], which states that the fundamental groups of cubical complexes of Example 3.3 embed in right-angled Artin groups. Or as well, the proof mimics that of [28], which is based on hyperplane properties.

Sketch of Proof: Given a finite graph Γ, let T_Γ denote the NPC cubed complex which is a finite Eilenberg-MacLane space for the Artin right-angled group with generators $V(\Gamma)$ and commutators given by edges in $E(\Gamma)$. In this space, there is an n-torus built from a cube with sides identified for each clique of n vertices in Γ (see [11] for details of the construction). Given a reconfigurable system, let Γ denote the graph whose vertex set is the set of generators for the system and whose edges connect vertices corresponding to local moves which commute at some state.

One constructs a well-defined cellular map from $\mathcal{S}$ to T_Γ by sending an n-cube of $\mathcal{S}$ to the n-cube torus in T_Γ defined by the n commuting generators defining the cube in $\mathcal{S}$. One checks easily (as in the proof of [11, Thm. 2]) that the link of each vertex maps injectively to a full subcomplex of the link of the target vertex. From Lemmas 1.41.6 of [9], the map from $\mathcal{S}$ to T_Γ is a local isometric embedding. A local isometric embedding into an NPC space is π_1-injective [6, Prop. 4.14]. □

Several algebraic results follow from this property. For example, fundamental groups of state complexes are linear, since right-angled Artin groups embed in $GL_n(\mathbb{R})$ [14].

Theorem 5.7 also follows easily from a more general result of Haglund and Wise [28], who prove the Artin right-angled embedding property for a related class of cube complexes:

Definition 5.8. A compact NPC cube complex X is said to be **A-special** if

(1) Each hyperplane in X embeds.
(2) Each hyperplane in X is 2-sided.
(3) No hyperplane in X directly self-osculates.
(4) No hyperplanes in X inter-osculate.

The meaning of (1) is that no hyperplane intersects any cube in more than one midplane. By (2) it is meant that the complement of a hyperplane disconnects its neighborhood. Condition (3) means that no two edges dual to a hyperplane can themselves intersect at a single vertex (with opposite orientation). Finally, (4) means that if two hyperplanes intersect, then any two edges dual to these hyperplanes which themselves intersect must span a 2-cell. See Fig. 11 for illustrations of obstructions to being A-special.

The result of [28] on fundamental groups of A-special complexes combines with the following observation to give an alternate proof of Theorem 5.7.

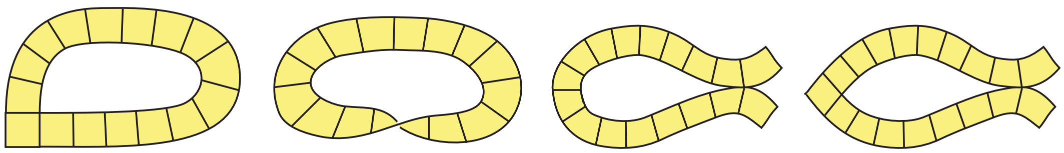

FIGURE 11. Four illegal subcomplexes of an A-special complex. From left-to-right: intersecting, 1-sided, self-osculating, and inter-osculating hyperplanes.

Proposition 5.9. *State complexes are A-special.*

PROOF: Properties (1), (2), and (3) are a direct consequence of Lemma 5.5. Property (4) is implied by the fact that commutativity of generators in a reconfigurable system is independent of the state at which they are applicable. Hence, if two generators commute at a given state, then they commutate at any other state to which both generators are applicable. □

The class of A-special cube complexes is, however, strictly larger than the class of state complexes (making the results of [28] stronger).

Proposition 5.10. *Not every A-special cube complex is a state complex.*

Proof. Consider the space given by sewing two squares together at their corners using a half twist in one. There are four states $\{\mathbf{u}_i\}_0^3$ and four distinct generators $\{\phi_0, \phi_1, \psi_0, \psi_1\}$, where the ϕ_i commute and the ψ_i commute, but no ϕ_i commutes with any ψ_i: see Fig. 12. One observes that this is A-special. At state $\mathbf{u}_0$ we have the following behavior:

$$\phi_1\phi_0[\mathbf{u}_0] = \psi_1[\mathbf{u}_0]. \tag{5.4}$$

Consider the set of vertices $\Delta_{0,3}$ on which states $\mathbf{u}_0$ and $\mathbf{u}_3$ differ.

$$\Delta_{0,3} = \{v \in V(\mathcal{G}) : \mathbf{u}_0(v) \neq \mathbf{u}_3(v)\} \tag{5.5}$$

Since $\psi_1[\mathbf{u}_0] = \mathbf{u}_3$, it follows that $\Delta_{0,3} \subset \text{TR}(\psi_1)$. However, as $\phi_0\phi_1[\mathbf{u}_0] = \mathbf{u}_3$ and the ϕ_i commute, it also follows that $\Delta_{0,3} \cap \text{TR}(\phi_i) \neq \emptyset$ for $i = 0, 1$. Moreover, since $\phi_0[\mathbf{u}_0] = \psi_0[\mathbf{u}_0]$, we have that $\Delta_{0,3} \cap \text{TR}(\psi_0) \neq \emptyset$. Together, these statements imply that $\text{TR}(\psi_0) \cap \text{TR}(\psi_1) \neq \emptyset$, contradicting the assumption that ψ_0 and ψ_1 commute. □

We note that what prevents this complex from being realizable as a state complex has nothing to do with the fact that multiple cells share the same vertex sets: a simple subdivision yields a nicer combinatorial structure but does not alter the above proof.

5.4. **Graph products.** A great many NPC cube complexes are indeed realizable as state complexes.

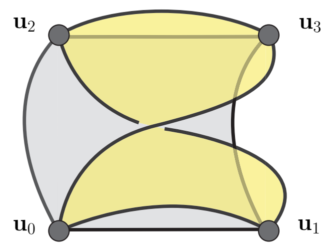

FIGURE 12. An example of an NPC cube complex which is A-special but not realizable as a state complex.

Theorem 5.11. *Any finite connected NPC subcomplex of a product of graphs can be realized as the state complex of a local reconfigurable system.*

PROOF: Suppose X is the subcomplex. As it is finite, we can regard X as a subcomplex of a finite product of finite graphs $\Gamma_1 \times \Gamma_2 \times \cdots \times \Gamma_N$. We set the domain $\mathcal{G} = \coprod \Gamma_i$ to be the disjoint union of the graphs. The alphabet is $\mathcal{A} = \{0, 1\}$. States for the system will consist of labelings which have a single '1' label in each Γ_i, all other vertices being labeled '0'. Loosely speaking, generators will correspond to sliding a '1' label along some edge of some Γ_i, as in Example 3.5.

For each hyperplane $\mathcal{H}$ of X we define a generator $\phi_{\mathcal{H}}$ as follows. Each edge in X transverse to $\mathcal{H}$ corresponds to a unique edge e in some Γ_i (see proof of Lemma 5.4). We set $\text{TR}(\phi_{\mathcal{H}}) = e$. The support of $\phi_{\mathcal{H}}$ is chosen to get the proper commutativity with other moves. Let $p_j : \prod_i \Gamma_i \to \Gamma_j$ denote projection of the direct product to the j^{th} factor. Define

$$\text{SUP}(\phi_{\mathcal{H}}) = \mathcal{G} - \bigcup_{j \neq i} p_j(\mathcal{H})$$

The local states have all vertices of SUP labeled '0' except for the single edge in TR, which has '0' and '1' (exchanged) at the boundary vertices.

Define a map $\Psi : X \to \mathcal{S}$ as follows. Each 0-cell $v \in X^{(0)}$ is an ordered N-tuple of vertices $v = (v_i) \in \prod_i \Gamma_i$. Define $\Psi(v)$ to be the state given by labeling each $v_i \in \Gamma_i$ with '1', all other vertices being labeled '0'. We extend Ψ as follows. Let C be an n-dimensional cube in X corresponding to the product of n edges $e_i \in \Gamma_{\alpha_i}$. Then there are exactly n hyperplanes $\{\mathcal{H}_i\}_1^n$ of X intersecting C, each corresponding to the edge e_i. The generators $\phi_{\mathcal{H}_i}$ are distinct and commute since

$$\text{SUP}(\phi_{\mathcal{H}_i}) \cap \text{TR}(\phi_{\mathcal{H}_j}) = \{e_j\} \cap \left(\mathcal{G} - \bigcup_{k \neq i} e_k \right) = \emptyset.$$

Define $\Psi(C)$ to be the n-cube in $\mathcal{S}$ defined by the commuting generators $(\phi_{\mathcal{H}_i})_1^n$. One easily verifies that Ψ is a bijection between cubes in X and cubes in $\mathcal{S}$ and furthermore that it respects the gluings between the cubes. Hence, Ψ gives the desired isometry between X and $\mathcal{S}$. □

It is an open question which manifolds can be represented as an NPC cube complex. The best converse to Theorem 5.11 would be that any state complex embeds as a subcomplex of a product of graphs. Unfortunately, this is not true.

Proposition 5.12. *The 2-d NPC complex of Fig. 13 is a realizable state complex which does not embed as a subcomplex of a product of graphs.*

PROOF: Let $\mathcal{G}$ be equal to the disjoint union of two closed edges and let $\mathcal{A} = \{0, 1\}$. Define three generators for the reconfigurable system. The first, ϕ_1, corresponds to exchanging $0 \leftrightarrow 1$ along the first edge in $\mathcal{G}$; the support and the trace are equal to this edge. Likewise, ϕ_2 with the second edge in $\mathcal{G}$. These two moves clearly commute. The third generator, ϕ_3, has support and trace equal to $\mathcal{G}$, and has the effect of performing both ϕ_1 and ϕ_2. This move commutes with no other generators.

Observe that any loop of three edges in the 1-skeleton of a graph product must lie entirely in one factor. This is because the projection of such a loop to each factor must also be a loop and therefore have at least two edges. The state complex of the above system contains a cycle of length three. As two of these edges commute, they must not lie in the same factor: contradiction. □

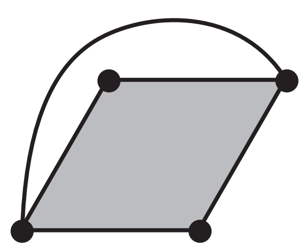

FIGURE 13. A state complex which cannot be a subcomplex of a graph product.

Note however that a suitable subdivision of the edge associated to ϕ_3 – in effect making it longer by replacing it with four generators – yields a homeomorphic state complex which does embed as a subcomplex of a unit cube. It may be the case that any realizable state complex has a subdivision which embeds as a subcomplex of a graph product.

Although this graph embedding property fails in general, the following partial solution shows that any embedding problems stem from noncontractible loops in the complex. The following simple result is likely well-known to experts: we include it for completeness.

Proposition 5.13. *Any finite CAT(0) cubical complex is a subcomplex of an N-cube for N sufficiently large.*

PROOF: Let X be a finite CAT(0) cubical complex and let N denote the total number of hyperplanes in X. Assume without a loss of generality that all edges in X are $[-1,1]$. We construct an embedding $\Psi : X \to [-1,1]^N$ as follows. For each hyperplane $\mathcal{H}$ and point p in X let $d(p,\mathcal{H})$ denote the geodesic distance from p to $\mathcal{H}$. Note that because X is CAT(0) , $d(p,\mathcal{H})$ is well-defined [6, Prop. 2.4]. Furthermore, since hyperplanes divide a CAT(0) cube complex into exactly two pieces [30], one can assign a transverse orientation to each $\mathcal{H}$. Define

$$\Psi_i(p) := \pm \frac{d(p,\mathcal{H}_i)}{\max\{1, d(p,\mathcal{H}_i)\}}$$

where the $\pm$ is assigned to be consistent with the chosen transverse orientation to the i^{th} hyperplane. The map Ψ is continuous since the geodesic distance $d(p,\mathcal{H}_i)$ is a continuous function in p and normalizing distances is distance non-increasing. Note that Ψ is an embedding from $X^{(0)} \to \{-1,1\}^N$ since each vertex of X is at least a unit distance from every midplane of every cube in X and distinct vertices of X are separated by some midplane and hence receive opposite signs under Ψ in that coordinate.

Assume inductively that Ψ embeds the $(k-1)$-dimensional skeleton $X^{(k-1)}$ to the $(k-1)$-dimensional skeleton of $[-1,1]^N$ and let C denote a k-dimensional cube in X. The boundary ∂C is sent by Ψ to the boundary of a k-dimensional face F of $[-1,1]^N$. Within the interior of C, any geodesic distance to any hyperplane which does not intersect C is greater than or equal to 1. Thus, the k midplanes of C completely determine the non-unit values of Ψ on the interior of C and $\Psi(C) = F$.

Should any two k-cubes be sent by Ψ to the same face F of $[-1,1]^N$, then Ψ^{-1} of the 'center' of F would consist of the mutual intersections of the k hyperplanes in these cubes: a pair of disjoint points, one in the center of each k-cube. As hyperplanes are totally geodesic subsets of a CAT(0) space [30], and, as intersections of totally geodesic subspaces are still totally geodesic, we have a contradiction.

Thus, Ψ is an embedding on $X^{(k)}$. This completes the induction step and the proof. □

Of course, this is not an isometric embedding.

Corollary 5.14. *Any finite CAT(0) cubical complex is realizable as the state complex of a reconfigurable system.*

In particular, the spaces of phylogenetic trees defined by Billera, Holmes, and Vogtmann [5] are state complexes of a reconfigurable system.

6. Questions

There are a number of open questions concerning the geometry, topology, and algebra of state complexes.

Question 6.1. Are state complexes really "discretizations" of some "continuous" configuration space? There are certain examples of reconfigurable systems for which it makes sense to refine the underlying graph and obtain a sequence of reconfigurable systems. In such examples, one can ask whether the sequence of state complexes enjoys any sort of convergence properties. A canonical example of such refinement occurs in the system of Example 3.3. Consider a refinement of the underlying lattice of Γ which inserts additional vertices with the zero label along edges.

It follows from the work of Abrams [1] that (in our terminology) the state complex of this refined system stabilizes in homotopy type: after a fixed number of refinements, all further refinements have homotopy equivalent state complexes. Furthermore, this "stabilized" state complex is in fact homotopic to the topological configuration space of N points of Γ, the N-fold product of the graph minus the pairwise diagonal.

One can certainly construct examples for which this type of refinement does not lead to state complexes which stabilize in homotopy type. However, it may be that there is a notion of refinement for which convergence in the Gromov-Hausdorff sense works. A first step is to formalize the notion of refinement and classify what types of convergence properties hold and when.

Question 6.2. What can be said about the homology of state complexes? It was shown in [22] that the configuration space of N points on a graph Γ has homological dimension bounded above by the number of essential vertices of Γ (vertices of degree greater than two), independent of N. Thanks to [1] this means that the state complex approximations from Example 3.3 have a bound on the homological dimension which is perhaps far below that of the topological dimension of the cube complex. This question is particularly interesting in combination with the previous question on refinement and convergence.

Question 6.3. To what extent is the nonpositive curvature present in state complexes prevalent in physical settings? In many fields of mathematics, one finds that there is a large, interesting subclass of objects whose natural underlying hyperbolic structure allows for good theorems. This is certainly true in dynamical systems [hyperbolic dynamics and the Smale program], 3-manifolds [the hyperbolic 3-manifolds being both interesting and prevalent], and group theory [Gromov-hyperbolic groups being both interesting and prevalent]. To what extent does this meta-principle hold in physical systems? For example, does the natural local CAT(0) geometry of the state complex in the protein folding system of Example 3.6 explain the *Levinthal paradox* — the observation that chains with enormously large configuration spaces relax to a stable conformation in an extremely short amount of time? A state complex with lots of negative curvature could explain such behavior, as the volume in a hyperbolic space is exponential in radius. A simple energy-gradient descent on such a hyperbolic configuration space could explain why protein chains do not 'waste time' in getting to their preferred conformations.

Question 6.4. Is there a better way to complete transition graphs to higher dimensional objects? We have used cubical complexes as a completion. In many respects, these are natural — many of the examples in Section 3 attest to this. Nevertheless, there are likely other ways to fill in the transition graph to get a cell complex with beneficial properties. Permutohedra and associahedra are examples of transition graphs for reconfigurable systems which are completed to polytopes. Likewise, the Cayley complex is a useful completion of a Cayley graph. We note also the **Hom complexes** of Lovász [29]: $Hom(H, G)$ is a polyhedral complex whose vertices are graph homomorphisms $H \to G$ and whose cells are functions from $V(H)$ to sets in $V(G)$. Related constructions, like box- and neighborhood- complexes, are not necessarily cubical complexes, but can give rise to interesting manifolds [12], and have proven efficacious in solving combinatorics problems.

Question 6.5. Though we have given several results on the realization problem for state complexes, a complete characterization remains unknown. Which groups arise as fundamental groups of state complexes? Can one characterize the complexes themselves? Which 3-manifolds, e.g, are realizable? Is there a finite set of operations on cube complexes which generate all state complexes? (E.g., state complexes are clearly closed under products — take the disjoint union of the reconfigurable systems.)

Acknowledgements

Many of the details in the definition of reconfigurable systems and state complexes were developed in conversations with A. Abrams. Abrams' thesis [1] is in many respects the inspiration for this work. The first author has benefitted greatly from conversations with D. Farley, I. Kapovich, L. Sabalka, and P. Schupp. We are also grateful to D. Wise for sharing a draft of [28]. Finally, an unknown referee helpfully pointed out several unclear and inaccurate points.

References

[1] A. Abrams. *Configuration spaces and braid groups of graphs.* Ph.D. thesis, UC Berkeley, 2000.
[2] A. Abrams and R. Ghrist. Finding topology in a factory: configuration spaces. *Amer. Math. Monthly*, 109 (2002) 140–150.
[3] A. Abrams and R. Ghrist. State complexes for metamorphic robot systems. *Intl. J. Robotics Research*, 23 (2004) 809–824.
[4] S. Alexander, R. Bishop, and R. Ghrist. Pursuit and evasion on non-convex domains of arbitrary dimensions. In Proc. Robotics: Systems and Science, 2006.
[5] L. Billera, S. Holmes, and K. Vogtmann. Geometry of the space of phylogenetic trees. *Adv. Applied Math.* 27 (2001) 733-767.
[6] M. Bridson and A. Haefliger. *Metric Spaces of Nonpositive Curvature*, Springer-Verlag, Berlin, 1999.

[7] Z. Butler, S. Byrnes, and D. Rus. Distributed motion planning for modular robots with unit-compressible modules. In Proc. IROS, 2001.
[8] Z. Butler, K. Kotay, D. Rus, and K. Tomita. Cellular automata for decentralized control of self-reconfigurable robots. In Proc. IEEE ICRA Workshop on Modular Robots, 2001.
[9] R. Charney. The Tits conjecture for locally reducible Artin groups. *Internat. J. Algebra Comput.* 10 (2000) 783-797.
[10] G. Chirikjian. Kinematics of a metamorphic robotic system. In Proc. IEEE ICRA, 1994.
[11] J. Crisp and B. Wiest. Embeddings of graph braid groups and surface groups in right-angled Artin groups and braids groups. *Alg. & Geom. Top.* 4 (2004) 439–472.
[12] P. Csorba and F. Lutz. Graph coloring manifolds. To appear, *Contemp. Math.* Preprint, 2005.
[13] M. Davis. Groups generated by reflections and aspherical manifolds not covered by Euclidean space. *Ann. Math.* (2) 117 (1983) 293–324.
[14] M. Davis and T. Januszkiewicz. Right angled Artin groups are commensurable with right angled Coxeter groups. *J. Pure and Appl. Algebra* 153 (2000) 229–235.
[15] D. Epstein et al. *Word Processing in Groups*. Jones & Bartlett Publishers, Boston MA, 1992.
[16] M. Erdmann and T. Lozano-Perez. On multiple moving objects. In Proc. IEEE ICRA, 1986.
[17] R. B. Fair, V. Srinivasan, H. Ren, P. Paik, V. K. Pamula, and M. G. Pollack. Electrowetting based on chip sample processing for integrated microfluidics. *IEEE Inter. Electron Devices Meeting (IEDM)* 2003.
[18] R. B. Fair, A. Khlystov, V. Srinivasan, V. K. Pamula, and K.N. Weaver. Integrated chemical/biochemical sample collection, pre-concentration, and analysis on a digital microfluidic lab-on-a-chip platform. In *Lab-on-a-Chip: Platforms, Devices, and Applications, Conf. 5591, SPIE Optics East*, Philadelphia, Oct. 25–28, 2004.
[19] D. Farley. Finiteness and CAT(0) properties of diagram groups. *Topology* 42 (2003) 1065–1082.
[20] P. Gaucher. About the globular homology of higher dimensional automata. *Cahiers de Top. et Geom. Diff. Categoriques.* 43(2) (2002) 107–156.
[21] R. Ghrist. Shape complexes for metamorphic robot systems. In *Algorithmic Foundations of Robotics V, STAR 7* (2004) 185–201.
[22] R. Ghrist. Configuration spaces and braid groups on graphs in robotics. In *AMS/IP Studies in Mathematics* 24 (2001) 29–40.
[23] R. Ghrist and D. Koditschek. Safe, cooperative robot dynamics on graphs. *SIAM J. Control Optim.* 40 (2002) 1556–1575.
[24] R. Ghrist and S. LaValle. Nonpositive curvature and Pareto optimal motion planning. To appear, *SIAM J. Control Optim.*, 2006.
[25] R. Ghrist, J. O'Kane, and S. M. LaValle. Computing Pareto optimal coordinations on roadmaps. *Intl. J. Robotics Research*, 12 (2006) 997–1012.
[26] M. Gromov. Hyperbolic groups. In *Essays in Group Theory*, MSRI Publ. 8, Springer-Verlag, 1987.
[27] E. Goubault. Schedulers as interpreters of higher-dimensional automata. In Proc. Workshop on Partial Evaluation and Semantic-Based Program Manipulation (PEPM), 1995.
[28] F. Haglund and D. Wise. Special cube complexes. Preprint, March 2005.
[29] L. Lovász. Kneser's conjecture, chromatic number, and homotopy. *J. Combin. Theory Ser. A*, 25 (1978) 319–324.
[30] G. A. Niblo and L. D. Reeves. The geometry of cube complexes and the complexity of their fundamental groups. *Topology*, 37 (1998) 621–633.
[31] V. Pratt. Modelling concurrency with geometry. In Proc. 18th Symp. on Principles of Programming Languages, 1991.
[32] M. Raussen. State spaces and dipaths up to dihomotopy. *Homotopy, Homology, & Appl.*, 5 (2003) 257–280.
[33] M. Raussen. On the classification of dipaths in geometric models of concurrency. *Math. Structures Comp. Sci.* 10 (2000) 427–457.
[34] L. D. Reeves. Rational subgroups of cubed 3-manifold groups. Michigan Math. J. 42 (1995) 109–126.
[35] C. Reutenauer. *The Mathematics of Petri Nets*. Prentice-Hall, 1990.

[36] A. Sali, E. Shakhnovich, and M. Karplus. How does a protien fold? *Nature* 369 (1994) 248–251.
[37] A. Sali, E. Shakhnovich, and M. Karplus. Kinetics of protien folding. *J. Mol. Biol.* 235 (1994) 1614–1636.
[38] J. Swiatkowski. Estimates for homological dimension of configuration spaces of graphs. *Colloq. Math.* 89 (2001) 69–79.
[39] J. Walter, J. Welch, and N. Amato. Distributed reconfiguration of metamorphic robot chains. In Proc. ACM Symp. on Distributed Computing, 2000.
[40] J. Walter, E. Tsai, and N. Amato. Choosing good paths for fast distributed reconfiguration of hexagonal metamorphic robots. In Proc. IEEE ICRA, 2002.
[41] J.E. Walter, J.L.Welch, and N.M. Amato. Concurrent metamorphosis of hexagonal robot chains into simple connected configurations. *IEEE Trans. Robotics & Automation* 15 (1999) 1035–1045.
[42] M. Yim, Y. Zhang, J. Lamping, and E. Mao. Distributed control for 3-d metamorphosis. *Autonomous Robots J.* 10 (2001) 41–56.
[43] E. Yoshida, S. Murata, K. Tomita, H. Kurokawa, and S. Kokaji. Distributed formation control of a modular mechanical system. In Proc. Intl. Conf. Intelligent Robots & Sys., 1997.

DEPARTMENT OF MATHEMATICS, COORDINATED SCIENCE LABORATORY, UNIVERSITY OF ILLINOIS, URBANA IL, 61801

E-mail address: ghrist@math.uiuc.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF ILLINOIS, URBANA IL, 61801

E-mail address: vpeterso@math.uiuc.edu