http://topology.auburn.edu/tp/

TOPOLOGY PROCEEDINGS

Volume 45, 2015

Pages 73–109

http://topology.nipissingu.ca/tp/

State Complexes and Special Cube Complexes

by

Valerie J. Peterson

Electronically published on April 5, 2014

Topology Proceedings

Web:	http://topology.auburn.edu/tp/
Mail:	Topology Proceedings
	Department of Mathematics & Statistics
	Auburn University, Alabama 36849, USA
E-mail:	topolog@auburn.edu
ISSN:	0146-4124

http://topology.auburn.edu/tp/
http://topology.nipissingu.ca/tp/
TOPOLOGY PROCEEDINGS
Volume 45 (2015)
Pages 73-109
E-Published on April 5, 2014

STATE COMPLEXES AND SPECIAL CUBE COMPLEXES

VALERIE J. PETERSON

Abstract. State complexes are cube complexes that record the discrete motions of a reconfigurable system: a physical or abstract setting in which local movements effect global changes in the shape of the system. In addition to possessing interesting geometric and topological features (e.g., they are non-positively curved, aspherical spaces), state complexes are also examples of special cube complexes and hence have linear fundamental groups. This paper presents new insights into state complexes as a subclass of special cube complexes, establishing clear contrasts between the two. Specific obstructions to realizing state complexes are presented and examples are chosen to illustrate that, unlike *special*, the property *state* is not necessarily inherited by subcomplexes or finite covers.

1. Introduction

A great many situations arise in the physical world in which some dynamically changing system – one involving multiple moving parts, for example – must be rearranged in a controlled manner. Motion-planning problems in robotics are one such setting: one wishes to find a method for (optimally) reconfiguring a collection of independent robotic agents within a shared workspace. Another is found in manufacturing, in which the success of an assembly task depends on executing a large number of individual movements. Still others exist in biology and chemistry, as well as in more abstract settings, such as group theory and computer science; myriad examples appear in [20].

2010 *Mathematics Subject Classification.* Primary 57Q05, 53C23, 20F65; Secondary 20F36, 20F67.

Key words and phrases. CAT(0) cube complexes, non-positive curvature, right-angled Artin groups, special cube complexes.

A classical approach to modeling these types of systems is to build a *configuration space*, whose points record the allowable states of the system; points corresponding to illegal states (e.g., those representing collisions between independent agents) are deleted. Studying the configuration space reveals information about the underlying system and how to rearrange its components effectively. While the viewpoint we adopt here is similarly topological in flavor, the version of a configuration space we consider is a "discretized" one, capturing movements in systems which themselves can be regarded as discrete. This is not as unnatural as it might sound. Indeed, there are a wide variety of situations that can be accurately described in terms of discrete motions; we will call such systems *reconfigurable.* Several examples of reconfigurable systems appear in the next section, and others can be found in [20].

We wish to find a useful geometric representation of a reconfigurable system, similar in spirit to the classical configuration space but of a more combinatorial nature, reflecting the inherent structure of the underlying discrete system. The intuitive starting point is to build a *transition graph* whose vertices correspond to states of the system and whose edges represent the elementary transitions ("moves") between states. This is analogous to the construction of a Cayley graph for a group presentation, except for the fact that the transition graph may not be homogeneous – not all "generators" may be applicable at a given state.

The notion of a transition graph is not new; it has been adapted for use by the robotics community [12] and in various social sciences [10], among others. In what follows, however, we extend the transition graph combinatorially by regarding it as the 1-skeleton of a higher dimensional cube complex, wherein a k-dimensional cube is present for each collection of k independent moves: this is the *state complex.* Each k-cube records the fact that is possible to perform the given tasks concurrently, thereby encoding information about optimizing reconfiguration. While increasing the dimension of the space and adding information in the form of cubical cells may seem to increase its complexity, the presence of higher dimensional cubes actually simplifies the space in important geometric, topological, and group theoretic ways and allows us to exploit the piecewise-Euclidean structure (using, e.g., Gromov's combinatorial "link condition"). The first property we find is *non-positive curvature*, detected with the use of versatile tools from CAT(0) geometry. This will imply that the spaces are aspherical, as they have CAT(0) (and hence contractible) universal covers. We also see that fundamental groups of such spaces are all subgroups of *right-angled Artin groups*, objects which have attracted much interest in the last several decades.

A fair amount was known about this inherent geometry and topology for a particular subset of state complexes, called *discretized configuration spaces*, owing to Abrams in [1] and Ghrist in [19], but here we consider the more general complexes, as developed in [2], [3], and [20]. In particular, we draw a number of connections between state complexes and related objects called *special cube complexes*. Special cube complexes were introduced by Haglund and Wise in [22] and, despite being produced by far more abstract motivations than the manufacturing or robotics settings mentioned above, are closely related to state complexes. This relationship ties practically-motivated state complexes to more theoretical pursuits.

This article undertakes a comparative study of state complexes and special cube complexes as an early step in classifying state complexes. Though a complete classification is certainly desirable it seems unlikely: we will see that any comprehensive description of state complexes must combine local requirements (not hard to describe) with global ones (much more difficult to characterize), and that current geometric, topological, and group theoretic tools are insufficient to distinguish state complexes from non-state complexes axiomatically. Hence, a large portion of the following paper is dedicated to presenting the subtleties of state complexes that have thus far obfuscated an entirely axiomatic description.

1.1. A motivating example.

Suppose you are a manufacturer wishing to coordinate robotic assembly agents moving within your factory. The robots travel on a shared track on the floor and must move between various workstations in order to accomplish a given task. Before beginning assembly, you need to know: *Can the task can be accomplished without collisions? If so, how might it be accomplished optimally?* The collection of robots on the factory floor is an example of a *reconfigurable system*: a system in which independent agents make small "local" moves that affect global positions of the system. The following canonical example nicely illustrates the premise underlying a reconfigurable system; it parallels an example of Abrams [1].

Example 1.1. *Robots on a track.* Elaborating on the factory setting suggested above, consider two robots moving on a track isomorphic to K_4, the complete connected graph on four vertices (Figure 1 (left)). Assume the robots have limited sensory capabilities and are able only to slide along an empty edge to an unoccupied vertex without stopping, backing up, or communicating with one another before or during movement. (This is often the case, as more sophisticated robots are cost-prohibitive.) In this context, movement along an edge is a "black box" and we may think of a robot hopping discretely from one vertex to an empty adjacent vertex.

Restricting our attention to discrete movements is in fact a natural perspective; readers are referred to [20] for many settings that are effectively described via this type of discrete motion.

As mentioned, a reasonable first step in recording the movements of the system is to build a *transition graph*, whose vertices represent configurations of the system, called *states*, and whose edges denote the moves between states. A robot moving from one vertex of K_4 to another yields two distinct states connected by an edge in the transition graph. This reduces the question of linking distinct configurations to finding a path between the appropriate vertices in the transition graph. Note that any such path avoids collisions by virtue of how we allowed the robots to move; points in the transition graph represent only legal configurations of robots on distinct vertices of K_4.

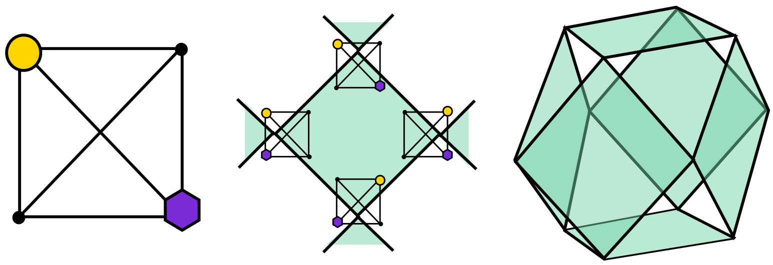

FIGURE 1. (*left*) Two robots move on the track K_4. (*center*) A local view of the state complex with states superimposed over vertices. (*right*) The state complex for two robots moving on K_4.

To address the practical concern of optimality in reconfiguration, independent moves should be parallelized when possible. We therefore add information to the transition graph in the form of higher dimensional cells – in this case, squares – whenever both robots can move simultaneously (i.e., along disjoint edges). The boundary of such a square denotes that moving Robot A and then Robot B is equivalent to moving B and then A; the diagonal of two-cell in the interior represents executing both moves concurrently. The result is the *state complex*, a cube complex which coordinates physically independent moves. A local picture of the state complex for this system of two robots is shown in Figure 1 (center) next to the entire state complex (right).

The next section formalizes these ideas.

2. Reconfigurable Systems and State Complexes

2.1. Reconfigurable systems.

Example 1.1 provides a good starting point which we now generalize. Begin by labeling two of the vertices of K_4 by 1 and 2 to denote the vertices on which the (distinct) robots sit, and label the empty vertices by 0. Now, states of the system are simply labelings of the vertices of K_4 by $\{0, 1, 2\}$ in which each of the labels 1 and 2 appears exactly once. Moves, which we call *generators*, exchange either adjacent 0 and 1 labels or adjacent 0 and 2 labels. This suggests the following definition.

Definition 2.1. (*Reconfigurable system*) A *reconfigurable system* consists of a *domain graph*, $\mathcal{G}$, a finite *alphabet* $\mathcal{A}$ of vertex labels, a collection of generators $\{\phi_i\}_{i\in I}$, and a collection of states, which is closed under the operation of applying generators. A *state* is a labeling of the vertices $V(\mathcal{G})$ by elements of $\mathcal{A}$; alphabet labels may be repeated but each vertex in $\mathcal{G}$ is assigned exactly one label in a given state. A *generator* ϕ is defined via the following three objects:

(i) the *support*, $\text{SUP}(\phi) \subset \mathcal{G}$, a subgraph of $\mathcal{G}$;
(ii) the *trace*, $\text{TR}(\phi) \subset \text{SUP}(\phi)$, a subgraph of $\text{SUP}(\phi)$;
(iii) 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.,

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

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

Intuitively speaking, the support of a generator is the amount of information needed to determine the legality of the move, whereas the trace is the precise subset of $\mathcal{G}$ on which ϕ changes vertex labels (i.e., where the move occurs physically). The trace and support for a given generator may coincide, as in Example 1.1 where each is a single edge and its vertices, but the support may also be a strictly larger subgraph of the domain graph than the trace. This is precisely what makes the definition of a reconfigurable system robust enough to capture the essence of a wide variety of situations (examples to follow). A further generalization from Example 1.1 is that generators in the system are allowed the flexibility to change many vertex labels at once as the trace is not restricted to a single vertex or two.

Definition 2.2. (*Admissibility*) 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

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

A reconfigurable system is said to be *locally finite* if the number of generators admissible at any state is finite, independent of state.

Remark 2.3. Since the local states of each generator are unordered (the local state labels '0' and '1' were assigned arbitrarily), it follows that any generator ϕ that is admissible at the state $\mathbf{u}$ is also admissible at the state $\phi[\mathbf{u}]$. This also implies that $\phi[\phi[\mathbf{u}]] = \mathbf{u}$.

2.2. The state complex.

We now construct the cube complex that captures all possible states of a reconfigurable system and the generators connecting them, extending the notion of a transition graph by regarding it as the one-dimensional skeleton of a larger cell complex.

Definition 2.4. (*Cube complex*) Let $I = [-1, 1] \subset \mathbb{R}$. A *Euclidean cube* is the product space I^k for some $0 \leq k$, conferred with the product metric. A *cube complex*, X, is a CW complex in which the k-cells are isometric to Euclidean cubes, I^k, and the attaching maps on the boundaries of k-cells restrict to isometries from each $(k-1)$-face in ∂I^k into X^{k-1}. The image of a k-cell of X is called a *k-cube*, though we will use the regular terminology *vertex, edge* and *square* for $k = 0, 1, 2$. Note that each k-cube in a cube complex is embedded.[1]

Definition 2.5. (*Commutivity*) A collection of generators $\{\phi_{\alpha_i}\}$ is said to *commute* if

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

Disjoint traces and supports for a given collection of generators indicates that these generators can be applied simultaneously without physical interference: the portion of $\mathcal{G}$ that one must look at to determine the legality of a move (the support) does not overlap with any of $\mathcal{G}$ on which the

[1]We distinguish here between the terms *cube*, *cubed*, and *cubical* complex, as these are sometimes used interchangeably in the literature. Here, *cubical complex* will mean the cubical equivalent of a simplicial complex, wherein any collection of 2^n vertices determines at most one cube. This is more stringent a requirement than for a *cube* complex (state complexes do not satisfy it in general). In contrast, a *cubed complex* is also a CW complex built from Euclidean cubes but without the requirement that k-cubes be embedded; they need only be immersed.

other moves actually occur (the remaining traces). We define the state complex to be the cube complex with a k-cube for each occurrence of k admissible commuting generators.

Definition 2.6. (*State complex*) The *state complex*, $\mathcal{S}$, of a reconfigurable system is the following abstract cube complex. Each k-cube of $\mathcal{S}$ is an equivalence class $[\mathbf{u}; (\phi_{\alpha_i})_{i=1}^k]$ where

(i) $(\phi_{\alpha_i})_{i=1}^k$ is a k-tuple of commuting generators;
(ii) $\mathbf{u}$ is some state for which all the generators $(\phi_{\alpha_i})_{i=1}^k$ are admissible; and
(iii) $[\mathbf{u}_0; (\phi_{\alpha_i})_{i=1}^k] = [\mathbf{u}_1; (\phi_{\beta_i})_{i=1}^k]$ if and only if the list (β_i) is a permutation of (α_i) and $\mathbf{u}_0 = \mathbf{u}_1$ on the set $\mathcal{G} - \bigcup_i \text{SUP}(\phi_{\alpha_i})$.

A state complex $\mathcal{S}$ is said to be *locally finite* if the underlying reconfigurable system for $\mathcal{S}$ is locally finite.

Note that state complexes are indeed cube complexes by our definition: the assumption that generators are nontrivial ensures that cubes in a state complex are embedded.

2.3. Selected examples.

We return to Example 1.1 to illustrate the precise definitions in this intuitive setting.

Example 2.7. *Robots on a track, revisited.* In the language above, the two robots moving on K_4 in Example 1.1 comprise a reconfigurable system with underlying graph $\mathcal{G} = K_4$ and alphabet $\mathcal{A} = \{0, 1, 2\}$ (different alphabet labels are denoted by vertices of different colors/shapes in Figure 1). There are two types of generators in the system, one for each robot, ϕ_1 and ϕ_2. The local states of ϕ_i evaluate to 0 at one of the endpoints of an edge and i at the other; thus, the support and trace for each generator are both equal to a single closed edge.

At any of the 12 states of the system, there are precisely two unoccupied vertices of K_4 to which either robot can move; this yields four edges incident to each of the 12 vertices in the state complex. In addition, on any given edge the move ϕ_1 commutes with exactly one occurrence of ϕ_2 (either along a parallel edge or opposite diagonal); this is true regardless of starting state, due to the symmetry of K_4. Thus, four edges and two squares meet at every vertex in $\mathcal{S}$, as shown in Figure 1 (center). The full complex $\mathcal{S}$ is shown on the right of the same figure.

If one repeats this example using K_5 as the underlying track for the two robots (see also [1, 20]), the state complex has a pleasing structure. The generators and alphabet remain the same as above. At any state of the system, there are precisely three unoccupied vertices of K_5 to which either robot can move; this yields six edges incident to every vertex in the state complex. In addition, on any given edge the move ϕ_1 now commutes with exactly two occurrences of ϕ_2, regardless of state. Every vertex in the state complex $\mathcal{S}$ therefore has a neighborhood with six 2-cells patched cyclically around it (Figure 2 (center)), so $\mathcal{S}$ is a closed cubical surface. One counts 20 vertices in $\mathcal{S}$, 60 edges, and 30 faces, giving an Euler characteristic for $\mathcal{S}$ of -10. A straightforward check reveals that the surface can be oriented; thus, the state complex is a cubical surface of genus six. The full complex $\mathcal{S}$ is shown in Figure 2 (right).

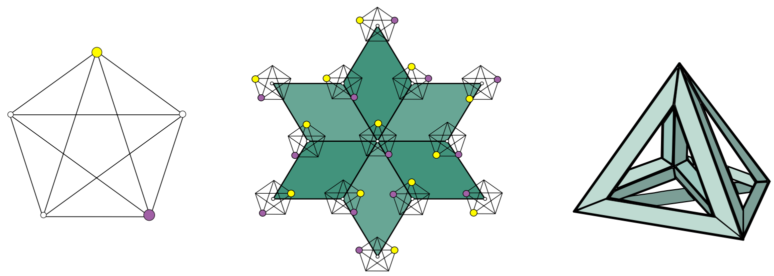

FIGURE 2. (*left*) Two robots move on the track K_5. (*center*) A local view of the state complex with states superimposed over vertices. (*right*) The entire state complex.

Example 2.8. *Molecular tilings.* In a 2008 experiment by Blunt et al. [7], scientists created a 2-dimensional network of p-terphenyl-$3, 5, 3', 5'$-tetracarboxylic acid (TPTC) molecules adsorbed on graphite that naturally exhibited a random tiling of the plane by rhombi, also called a *lozenge tiling*. The tiling they observed is non-periodic and non-homogeneous in the sense that at each vertex of the tiling, either 3, 4, 5, or 6 rhombic tiles can join together (Figure 3, center). Each rhombus in the tiling represents a single TPTC molecule; the various arrangements of tiles correspond to the different possible carboxylic–carboxylic hydrogen bonds that can form between molecules.

In addition to molecular tiles, the tiling contains what the authors refer to as "topological defects" in the form of empty triangular voids. The voids were observed to migrate through the network over time, forcing local rearrangements of tiles and giving rise to quasi-degenerate local minima within an energy landscape [7]. Voids in a tiling are illustrated in Figures 3 and 4 as unshaded triangles.

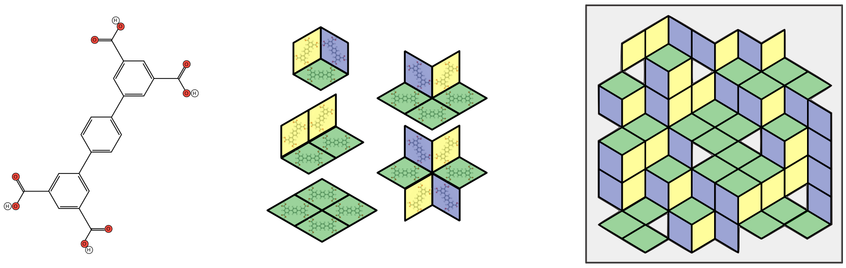

FIGURE 3. (*left*) A single TPTC molecule. (*center*) Molecules join together in arrangements of 3, 4, 5, or 6 rhombi. (*right*) A random tiling of the plane by TPTC molecules. Empty triangles allow molecules to move locally, transitioning the network between states of local energy minima.

We describe a lozenge tiling as a reconfigurable system by setting $\mathcal{G}$ equal to the triangular lattice dual to the tiling; we regard each rhombic tile as being composed of two equilateral triangles, and empty triangular voids are assigned vertices in this graph as well as rhombi. Vertices of $\mathcal{G}$ are labeled with the alphabet $\{0, 1, 2, 3\}$ according to empty or occupied vertices; each non-zero number in the alphabet represents one of three distinct tile orientations. The propagation of triangular defects is captured by generators that pivot a rhombic tile into an adjacent triangular void. There are three generators associated to every void, one for each free side of the triangle. Figure 4 shows the three directions in which tiles can pivot.

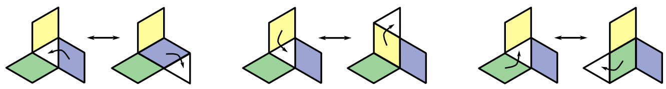

FIGURE 4. The three generators associated to any triangular void pivot rhombic tiles to fill the void from each open direction.

The trace of a generator contains the vertices associated to the rhombic tile it moves (two vertices) and the void it fills (one vertex). No two generators associated with the same void can commute, so the support of a generator must contain all vertices corresponding to neighbors of both the moving tile and the empty triangle. Note this is distinctly different than in the previous example; here, the support of each generator strictly contains the trace so as to accurately reflect the physical interdependence of moves. Generators affiliated with different voids, however, all commute (provided voids are sufficiently separated) so the upper bound on the dimension of cubes in the state complex depends on the total number of voids in the tiling.

This reconfigurable system is reminiscent of a metamorphic robotic system pioneered by Chirikjian [12]. The system consists of a (finite) collection of planar hexagonal robotic linkages locked together in an aggregate, forming a hexagonal lattice. Individual robotic agents may detach from their neighbors in the lattice and pivot around a fixed corner, similar to the movements of rhombic molecular tiles. Details for the reconfigurable system describing the hex-lattice metamorphic robots appear in [20].

Example 2.9. *Abstract cube complex.* When possible, realizing a specific cube complex as a state complex almost always requires defining the underlying reconfigurable system that generates it (though there do exist more general techniques, which appear in the next section). For the complex $\widehat{X}$ pictured below, we define this system abstractly as follows.[2]

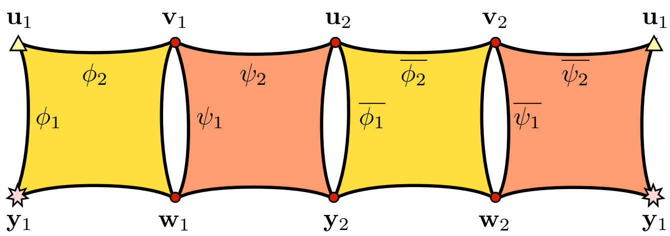

FIGURE 5. Generators and states for a reconfigurable system are superimposed on the complex, $\widehat{X}$.

We let

$$\mathcal{G} = \underset{a}{\bullet} \quad \underset{b}{\bullet} \quad \underset{c}{\bullet} \quad \underset{d}{\bullet} \quad \underset{e}{\bullet} \qquad \text{and} \qquad \mathcal{A} = \mathbb{Z}_2.$$

Generators for the system are defined in the table below. States and local states are of the form $\underline{\mathrm{a}}\,\underline{\mathrm{b}}\,\underline{\mathrm{c}}\,\underline{\mathrm{d}}\,\underline{\mathrm{e}}$, where empty slots in local states denote vertices of $\mathcal{G}$ that do not belong to the support of a given generator (and hence do not change).

[2]We will see later on that $\widehat{X}$ is actually a two-sheeted cover of a complex, X, that cannot be realized as a state complex for any reconfigurable system; this motivated the naming.

GENERATOR	TR	SUP	LOCAL STATES
ϕ_1	a	a, d, e	0 _ _ 0 0 ↔ 1 _ _ 0 0
ϕ_2	b, c	b, c, d, e	_ 0 0 0 0 ↔ _ 1 1 0 0
ψ_1	a	a, b, e	0 1 _ _ 0 ↔ 1 1 _ _ 0
ψ_2	c, d	b, c, d, e	_ 1 1 0 0 ↔ _ 1 0 1 0
$\overline{\phi_1}$	a	a, b, c	0 1 0 _ _ ↔ 1 1 0 _ _
$\overline{\phi_2}$	d, e	b, c, d, e	_ 1 0 1 0 ↔ _ 1 0 0 1
$\overline{\psi_1}$	a	a, c, d	0 _ 0 0 _ ↔ 1 _ 0 0 _
$\overline{\psi_2}$	b, e	b, c, d, e	_ 1 0 0 1 ↔ _ 0 0 0 0

States for the system are

$$\mathbf{u}_1 = 00000, \quad \mathbf{v}_1 = 01100, \quad \mathbf{u}_2 = 01010, \quad \mathbf{v}_2 = 01001,$$
$$\mathbf{y}_1 = 10000, \quad \mathbf{w}_1 = 11100, \quad \mathbf{y}_2 = 11010, \quad \mathbf{w}_2 = 11001.$$

The trace, support, and local states for each generator are chosen to reflect exactly the commutativity and admissibility shown in Figure 5. For example, ϕ_1 is admissible only at local states of the form 0 _ _ 0 0 and 1 _ _ 0 0 and therefore can be applied precisely at states $\mathbf{u}_1, \mathbf{v}_1, \mathbf{y}_1$ and $\mathbf{w}_1$. Further, ϕ_1 cannot commute with any generator whose support contains $\text{TR}(\phi_1) = \text{a}$, or any generator whose trace overlaps with $\text{SUP}(\phi_1) = \{\text{a, d, e}\}$, leaving only ϕ_2. Readers are encouraged to verify the proper commutativity is achieved for the remaining generators. Note as well that this choice of reconfigurable system generating $\widehat{X}$ is not unique.

Readers are referred to [20] and [28] for a bevy of further examples, including settings from biology (protein folding) [29], chemistry (digital microfluidics) [18], psychology (media theory and learning spaces) [10], robotics (metamorphic robotic systems) [3], combinatorics (permutohedra), and yet other fields that all admit descriptions as reconfigurable systems and thus give rise to state complexes.

3. Fundamental Properties and Realization of State Complexes

One of the most fascinating qualities of state complexes is the fact that they come equipped with an intrinsic metric that is non-positively curved, despite being composed of flat Euclidean cubes. Additionally, they are all examples of special cube complexes. With the intent of being self-contained, a fair amount of relevant background on state complexes is included here.

3.1. Curvature for cube complexes.

Much of our contemporary understanding of curvature for general metric spaces owes to fundamental work by A. D. Alexandrov [5], who defined what it means for a metric space X to have curvature bounded above by a real number κ by comparing triangles in X to those in a model space.[3] Later, Gromov famously dubbed this formulation the "CAT(κ) inequality" in honor of Alexandrov, along with Cartan and Toponogov, who had also made major contributions in the area [9]. In the case that X satisfies the CAT(κ) inequality for $\kappa = 0$, we simply refer to X as a CAT(0) space. A metric space is *non-positively curved (NPC)* if it satisfies the CAT(0) inequality only locally (i.e., for sufficiently small comparison triangles). Gromov is also responsible for reformulating these ideas in the context of cube complexes, providing purely combinatorial means to detect non-positive curvature in this setting [21].

Definition 3.1. (*Link, flag complex*) Let X be a simple[4] cube complex and let v be a vertex in X. The *link of* v, denoted $\mathcal{L}k[v]$, is the abstract simplicial complex with a k-cell for every $(k+1)$-cube in X incident to v. The boundary relations for k-simplices in the link are inherited from the boundary relations among the corresponding $(k+1)$-cubes in X. A simplicial complex K is a *flag complex* if any collection of vertices in K that are pairwise connected also span a simplex in K. Said another way, K is maximal among all simplicial complexes with the same 1-skeleton.

The link of every vertex in a state complex is simplicial by virtue of the following. First, all cubes are embedded, hence there can be no loops in the link of a vertex. Second, no two distinct squares may be glued along two adjacent edges, as these shared edges would represent generators producing the same cubical equivalence class, ensuring the two squares are in fact identified. Therefore, there can be no digons in the link (a *digon* is a pair of vertices connected by two distinct edges). A theorem of Gromov then asserts that global topological features of a cube complex are in fact determined entirely by the local behavior at vertex links: a finite dimensional Euclidean cube complex is non-positively curved if and only if the link of each vertex is a flag complex [21].

---

[3] The model spaces are: real hyperbolic space $\mathbb{H}^2$ with the distance function scaled by a factor of $\frac{1}{\sqrt{-\kappa}}$ for $\kappa < 0$, the Euclidean plane $\mathbb{E}^2$ for $\kappa = 0$, and the 2-sphere $\mathbb{S}^2$ with the distance function scaled by $\frac{1}{\sqrt{\kappa}}$ for $\kappa > 0$.

[4] The term *simple* guarantees that links are in fact simplicial. In particular, this prohibits two squares from being identified along two consecutive edges, as it would imply a non-simplicial digon in the link of their shared vertex.

As detailed above, links in a state complex are necessarily simplicial (so state complexes are simple cube complexes). Moreover, links are flag by virtue of Definition 2.5: whenever k generators commute pairwise at state $\mathbf{v}$, the entire collection must commute at $\mathbf{v}$. Translating this definition to vertex links, each pair of commuting generators gives rise to an edge in the link of $\mathbf{v}$, but the k-cube present at $\mathbf{v}$ (guaranteed by the commutativity of the entire collection of generators) ensures that the corresponding $(k-1)$-simplex in the link is filled in. Together with Gromov's link condition this shows the following.

Theorem 3.2. [20] *The state complex for any reconfigurable system is NPC.*

Non-positive curvature has both theoretical and practical implications. One useful and well-known property of (connected) CAT(0) spaces is that any two points are joined by a unique geodesic, which implies these spaces are contractible. Additionally, a complete geodesic metric space X is CAT(0) if and only if it is NPC and simply connected. ([9], for example, contains a thorough treatment of non-positively curved spaces.) The local property of non-positive curvature is nearly as nice: the Hadamard-Cartan Theorem implies that an NPC space X has a CAT(0) universal cover, so while X may itself allow multiple geodesics between two points, there is exactly one (local) geodesic per homotopy class. Practically speaking, the existence of a unique shortest path per homotopy class suggests how to address the motivating concern of optimality: given a path avoiding fixed obstacles in the workspace, there is an optimal way to reconfigure the factory robots. Additionally, if X is NPC, the contractibility of $\widetilde{X}$ ensures that all higher homotopy groups of X vanish (by Whitehead's Theorem), making X an Eilenberg-MacLane space of type $K(\pi_1, 1)$, also known as an *aspherical* space. An aspherical space X is determined up to homotopy entirely by its fundamental group, which is torsion-free whenever X is finite-dimensional [24].

3.2. Realizing state complexes.

As a first step in the classification of state complexes, we present what is currently known about recognizing arbitrary NPC cube complexes as state. Readers are referred to [20] for a thorough treatment of these results. Much of what follows makes use of the hyperplanes present in cube complexes, as developed in [27].

Definition 3.3. (*Hyperplane.*) Let X be a cubical complex, so that each cube in X is 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 $\{x_i = 0\}$ for some i. Said simply, a *hyperplane* is a union of midplanes glued together according to (restrictions of) the gluing maps between cubes.

More carefully, two midplanes M and N in X are *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 in some cube in X for every $i = 1, 2, \ldots, (n-1)$. A *hyperplane* is an equivalence class of midplanes with respect to this hyperplane equivalence; we generally denote a hyperplane by $\mathcal{H}$. An edge in a cube is *dual* to a midplane (and therefore a hyperplane) if it is perpendicular to the midplane.

Every edge in a state complex $\mathcal{S}$ represents the action of some admissible generator; edges that are parallel across squares, cubes, *etc.*, are dual to equivalent midplanes and therefore a common hyperplane. It is then unsurprising that each hyperplane $\mathcal{H}$ in a state complex $\mathcal{S}$ corresponds to the action of a unique generator, as in [20], Lemma 5.4. Further, each hyperplane belongs to a complex isomorphic to $\mathcal{H} \times [-1, 1]$ (Lemma 5.5). In particular, these combine to imply that the square complex shown in Figure 6 is not realizable as a state complex for any reconfigurable system, despite being non-positively curved. This provides us with a first non-example of state complex, demonstrating that the class of state complexes is a strict subset of the class of NPC cube complexes.

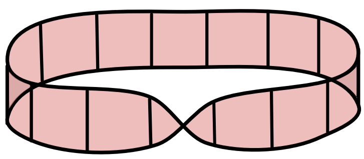

FIGURE 6. An NPC cube complex that is not a state complex.

On the other hand, cube complexes that can always be described as state complexes include appropriate subcomplexes of well-behaved products of graphs. In the following, a *simple graph* is a graph with no single-edge loops.

Theorem 3.4 ([20]). *Any finite NPC subcomplex of a product of simple graphs can be realized as a state complex for some reconfigurable system.*

The proof defines a reconfigurable system that will generate the given subcomplex, emulating the "robots on a graph" setting in Example 1.1. In short, the underlying graph for the system consists of the disjoint union of the graphs that appear as factors of the product, the alphabet is $\{0, 1\}$, and generators correspond to sliding robots along the edges of the individual graphs that appear in a given cube in the product, with traces and supports chosen to achieve the appropriate commutativity. The requirement "simple" ensures that no generators are trivial, leaving and returning to the same state.

For additional examples, note that since graphs are themselves trivially NPC and finite products of NPC spaces are again NPC [9], any finite product of simple graphs is itself automatically a state complex. Moreover, it is known that any finite CAT(0) cubical complex X can be embedded as a subcomplex of a n-dimensional cube, n being the number of hyperplanes in X [20], hence all finite CAT(0) cube complexes are also realizable as state complexes.

Beyond the somewhat rigid confines of graph products, one may also construct a number of examples of state complexes with homogeneous cubical structure. Paralleling a result of Davis [13], given any finite simplicial flag complex L one may construct a reconfigurable system whose state complex $\mathcal{S}$ satisfies $\mathcal{L}k[v] = L$ for all vertices $v \in \mathcal{S}$ [20]. It is therefore possible to construct an n-manifold state complex somewhat generically by choosing L to be a simplicial flag $(n-1)$-sphere, implying that a great many cubical surfaces and three-manifolds arise as state complexes. As an example, by choosing L to be a simplicial cycle of length four, the construction in [20] yields a two-torus state complex composed of 16 squares.

The best possible converse to Theorem 3.4 would be that any state complex embeds as a subcomplex of a product of graphs. This is unfortunately not true: the following complex can be given a state structure (the associated reconfigurable system has three generators) but cannot be embedded in a product of graphs.

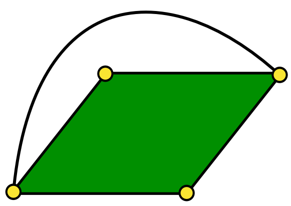

FIGURE 7. A state complex that does not embed in a product of graphs.

3.3. Special cube complexes.

We now consider cube complexes known as *A-special* and the related class of *special cube complexes*. Introduced by Haglund and Wise in [22], A-special (and special) complexes were developed as an attempt to generalize square complexes called *clean $\mathcal{VH}$-complexes* [31], [32] that exhibited a convenient group theoretic property (*canonical completion and retraction*) which we invoke a bit later. Beyond the original intent, special cube complexes quickly proved to be powerful tools in the study of finiteness properties (subgroup separability, e.g., [23]) and infinite groups, as well as fundamentally important in the topology of three-manifolds, featuring prominently in Agol's 2012 proof of the virtual Haken conjecture [4]. Special cube complexes are defined in terms of their *hyperplanes* (Definition 3.3) and their *walls*, which are dual to hyperplanes.

Definition 3.5. *(Wall)* A *wall* in a cube complex X is the set of edges dual to some fixed hyperplane, and is denoted $W(\mathcal{H})$.

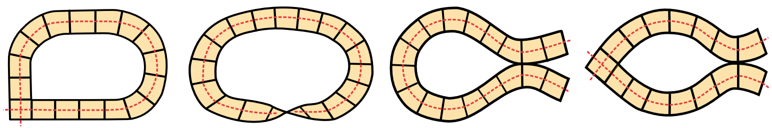

FIGURE 8. *(left to right)* Self-intersecting, one-sided, self-osculating, and inter-osculating hyperplanes are forbidden in an A-special cube complex.

Definition 3.6. *(Hyperplane pathologies)* The following hyperplane behaviors are depicted in Figure 8, from left to right: a *self-intersecting* hyperplane, a *one-sided* hyperplane, a *self-osculating* hyperplane, and a pair of *inter-osculating* hyperplanes.

(1) A hyperplane $\mathcal{H}$ in X *self-intersects* if it contains more than one midplane from the same cube. Equivalently, $\mathcal{H}$ self-intersects if the map $\mathcal{H} \to X$ is not injective, or if $\mathcal{H}$ has two dual edges that are also consecutive in some square. A hyperplane is *embedded* if it does not self-intersect.
(2) A hyperplane is *one-sided* if $W(\mathcal{H})$ cannot be given a consistent orientation. (By *consistent* we mean that whenever two edges appear opposite sides of a square they have the same orientation.) Equivalently, a hyperplane $\mathcal{H}$ is one-sided if its complement does not disconnect a neighborhood of $\mathcal{H}$. A hyperplane that is not one-sided is called *two-sided*.
(3) A hyperplane *self-osculates* if there exist two edges dual to the hyperplane, $a, b \in W(\mathcal{H})$, that share a vertex v but are not consecutive edges of any square containing v.
(4) Finally, a pair of hyperplanes *inter-osculates* if they both intersect and osculate.

Definition 3.7. *(Special cube complex)* Let X be a simple cube complex. X is *A-special* if it avoids all hyperplane pathologies listed in Definition 3.6. X is *special* if it has some finite cover that is A-special.

Clearly every A-special complex is also special. The same definitions apply if we replace "cube" with "cubed."

Example 3.8. Note that any graph is A-special. Further, any CAT(0) cube complex X is A-special. This follows from several standard facts about the hyperplanes of CAT(0) cube complexes [9], [27] that ensure none of the pathologies described in Definition 3.6 can occur. Finally, below we construct a cubed complex called the *Artin complex* that can be associated to a group (or to a cube complex). When associated with an appropriate group, this Artin complex is A-special.

Here is the motivation for introducing special cube complexes: all state complexes avoid the hyperplane pathologies listed above.

Theorem 3.9 ([20])**.** *State complexes are A-special cube complexes.*

The reasoning is straightforward: Pathologies (1), (2), and (3) in Definition 3.7 are all avoided in a state complex by virtue of Lemmas 5.4 and 5.5 of [20]. Pathology (4) is avoided due to the fact that commutativity of generators in a state complex is independent of their location. In particular, if two generators commute at a given state (indicated by the intersection of their hyperplanes), they commute everywhere they are admissible.

As readers may have discerned, the "A" in A-special refers to Artin, and specifically to a *right-angled Artin group*. General Artin groups arose as a generalization of braid groups and thus admit presentations similar to classical braid groups. Introduced by Baudisch in [6] and developed further over the next decade by Droms [15], [16], [17], Artin groups have attracted much attention since then, in part because of their actions on CAT(0) cube complexes. The focus here will be on *right-angled Artin groups*; these are groups for which all relators are commutators in specified pairs of generators. They are also referred to as *graph groups* as they may be put into correspondence with simplicial graphs.

Definition 3.10. *(Right-angled Artin group)* Let Γ be a simplicial graph, let $V(\Gamma)$ denote the vertices of Γ, let $E(\Gamma)$ denote the (geometric) edges of Γ. The vertices of Γ will give rise to generators of the group and the edges of Γ specify when generators commute. Specifically, the *right-angled Artin group* or *graph group* associated to Γ is the group presented by

$$\mathrm{A}(\Gamma) = \langle\, x_i,\ i \in V(\Gamma) \ \mid\ [x_i, x_j] = 1 \text{ for every } (i,j) \in E(\Gamma)\,\rangle. \tag{3.1}$$

From the definition, we see that free groups arise from graphs with no edges, and free abelian groups arise from complete graphs. Much is known about these groups (also known as RAAGs): most notably, right-angled Artin groups are linear; this was proved by Hsu and Wise in [25].

In addition, Davis and Januszkiewicz give a straightforward description of how to realize any finitely generated right-angled Artin group as a subgroup of a right-angled Coxeter group in [14], from which it follows that it is a subgroup of $SL_n(\mathbb{Z})$ for some n. Thus, embedding a group as a subgroup of a right-angled Artin group is one method for showing the group is linear. Defining a local isometry from a state complex into an Artin complex will induce such an embedding on the level of fundamental groups. We first develop some necessary notions of combinatorial maps.

Definition 3.11. *(Combinatorial map, local isometry)* A *combinatorial map* $f : X \to Y$ between CW complexes is one that sends open cells of X homeomorphically onto open cells of Y. If X and Y are cube complexes, a combinatorial map preserves k-cubes. A *combinatorial embedding* is a combinatorial map that is also a homeomorphism onto its image. Now assume that $f : X \to Y$ is a combinatorial map of simple cube complexes. The map f is an *immersion* if the induced map $\mathcal{L}k[v] \to \mathcal{L}k[f(v)]$ is an embedding for all $v \in X^{(0)}$. The map f is a *local isometry* if it is an immersion and $f(\mathcal{L}k[v])$ is a full subcomplex of $\mathcal{L}k[w]$. A subcomplex $A \subset B$ of a simplicial complex is said to be *full* if any simplex of B whose vertices lie in A is in fact entirely contained in A.

3.4. The Artin complex.

Definition 3.12. *(Artin complex)* Let $A(\Gamma)$ be a right-angled Artin group, as presented in Equation 3.1. Let X be the CW complex consisting of one vertex, an (arbitrarily oriented) edge loop for every generator x_i of A(Γ), and a 2-cell for each commutator in the set of relations for A(Γ). Each 2-cell is attached by labeling its boundary with the relator $x_i x_j x_i^{-1} x_j^{-1}$ and then gluing this boundary to the appropriately labeled loops. X is called the *standard 2-complex* for A(Γ). Note that X is a cubed complex whose fundamental group is $A(\Gamma)$.

We extend X to an NPC cubed complex by adding an n-cube in the form of an n-torus for each distinct collection of n pairwise commuting generators in A(Γ). The faces of such an n-torus are the 2-cells whose boundaries are labeled by commuting pairs of generators. This new complex is called the *Artin complex* associated to A(Γ) and is denoted $\mathcal{A}\text{RT}(\Gamma)$, adopting the terminology and notation used in [22].[5]

By virtue of its construction, an Artin complex, $\mathcal{A}\text{RT}(\Gamma)$, is an A-special cubed complex; details appear in [22].

[5]Readers may know this complex as the *Salvetti complex* associated to A(Γ), though this is somewhat of an abuse. The term *Salvetti complex* was originally used to denote the universal cover of this complex, which is a CAT(0) cube complex. The analogue of this complex in the case of spherical (i.e., finite type) Artin groups was introduced by Salvetti in [30].

Extrapolating, one may associate an Artin complex to any simple cube complex, X. Suppose first that all hyperplanes of X are embedded. To X we associate a simplicial graph Γ_X whose vertices correspond to hyperplanes $\mathcal{H}$ of X and whose edges connect pairs of intersecting hyperplanes; Γ_X is simplicial precisely because hyperplanes in X do not self-intersect. Now, to the graph Γ_X we associate an Artin complex. The procedure for building this complex is identical to the procedure we followed in Definition 3.12, except we proceed directly from the graph. Begin with a single vertex v, and attach a loop to v for every vertex in Γ_X. Since vertices of Γ_X correspond to hyperplanes of X, each loop is labeled by some hyperplane $\mathcal{H}_i \subset X$. To this 1-skeleton, attach a square for every edge in Γ_X. As each edge in Γ_X corresponds to a pair of intersecting hyperplanes, $\mathcal{H}_i \cap \mathcal{H}_j$, the boundary of this square is glued along loops $\mathcal{H}_i\mathcal{H}_j\mathcal{H}_i^{-1}\mathcal{H}_j^{-1}$. Similarly, glue in the boundary of a k-torus to the appropriate $(k-1)$-cells for every collection of k pairwise connected vertices in Γ_X. The resulting complex, $\mathcal{A}\textsc{rt}(\Gamma_X)$, is a non-positively curved cubed complex associated to the cube complex X via the graph Γ_X.

If, in addition to being embedded, the hyperplanes of X are all two-sided, the walls in X are therefore orientable which implies there is a combinatorial map τ sending X to $\mathcal{A}\textsc{rt}(\Gamma_X)$, as follows. τ sends all vertices of X to the single vertex $v \in \mathcal{A}\textsc{rt}(\Gamma_X)$, and all edges in the wall $W(\mathcal{H})$ to the loop labeled by $\mathcal{H}$, preserving orientation. Any k-cube $\sigma \in X$ is naturally mapped to the k-torus in $\mathcal{A}\textsc{rt}(\Gamma_X)$ arising from the k intersecting hyperplanes in σ. When X is an A-special cube complex, X and $\mathcal{A}\textsc{rt}(\Gamma_X)$ are nicely compatible; the following appears in [22].

Theorem 3.13 ([22]). *If X is A-special, the map $\tau : X \to \mathcal{A}\textsc{rt}(\Gamma_X)$ is a local isometry.*[6]

As argued above, the fact that hyperplanes in X are 2-sided and embedded guarantee the existence of the map τ, and the remaining two features of A-special complexes show that the map is indeed a local isometry. This is a crucial fact used in [22] to prove the following: If X is a compact, connected, special cube complex then $\pi_1 X$ embeds in a right-angled Artin group, and is therefore linear. This also implies that $\pi_1 X$ is residually finite, since finitely generated linear groups are residually finite [26].

In summary, we have thus far established the following facts about state complexes:

[6]In fact, the converse to the above statement is also true: a cube complex X is A-special if and only if there exists a graph Γ_X and a map $X \to \mathcal{A}\textsc{rt}(\Gamma_X)$ which is a local isometry.

- State complexes are NPC cube complexes and therefore aspherical.
- Finite NPC subcomplexes of graph products can be realized as state complexes, but not all state complexes embed in graph products.
- All CAT(0) cube complexes are state complexes.
- State complexes are A-special, and hence special.
- Fundamental groups of state complexes are subgroups of right-angled Artin groups, whose generators correspond to generators of the associated reconfigurable system and whose relations correspond to commuting generators. These groups are therefore linear and residually finite.

Though we have collected a variety of realization facts and properties, a complete axiomatic description of state complexes remains unrealized. This is because the underlying structure of a state complex is based on both local and global information: the individual interactions of commuting generators determine local states, but the overall connectivity of states is shaped by repeated application of these local generators and the resulting (global) hyperplane interactions, which is extremely hard to characterize. We therefore turn to an examination of the behavior of state complexes in relation to the larger set of special cube complexes in an attempt to add to our classification results.

4. State Versus Special

We have seen that every state complex for a locally reconfigurable system is actually an A-special cube complex. Here, we investigate ways in which the converse to this statement fails, beginning with three specific cases.

4.1. Special complexes that are not state.

At present, the only way to guarantee (abstractly) that an arbitrary non-positively curved A-special cube complex is a state complex is to embed it in a product of graphs, via Theorem 3.4 (though we have seen that failure to embed in a graph product is not necessarily an obstruction to being state). In general, recognizing a state complex requires defining the underlying reconfigurable system that generates it, so demonstrating that a complex *cannot* be state requires showing no such system can exist.

Theorem 4.1. *The A-special cube complexes pictured in Figure 9 cannot be realized as state complexes for any reconfigurable system.*

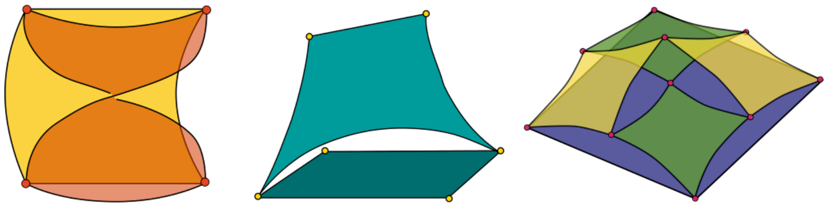

X Y Z

FIGURE 9. Non-positively curved A-special cube complexes that are not state.

Remark 4.2. Each of these NPC square complexes fails to admit a description as a state complex in a slightly different way. As we will see, the geometry of the complex X (Figure 9, left) prohibits generators from commuting as they should. The complex Y (center) is incomplete in the sense that edges (or squares) must be added to the complex in order to impose a state structure. The last complex, Z (right) is simply inconsistent with any possible system of state labels and generators one imposes. We address each complex in turn.

PROOF *(X is not state)*: This portion of the proof has been reproduced from [20] for completeness. We let X denote the space of "twisted squares" formed by gluing the four vertices of a square to the vertices of another, where a half-twist has been introduced in one square. That X is A-special is apparent from the simplicity of its hyperplanes, which are all single midplanes. Suppose X is the state complex for some reconfigurable system with underlying graph $\mathcal{G}$. Then there are four states in X, $\{\mathbf{u}, \mathbf{v}, \mathbf{w}, \mathbf{y}\}$, and four generators, $\{\phi_1, \phi_2, \psi_1, \psi_2\}$, one for each hyperplane of X (see Figure 10). By virtue of the squares present in the complex, ϕ_1 and ϕ_2 must commute, as must ψ_1 and ψ_2, but no ϕ_i commutes with any ψ_j.

Examining the generators at state $\mathbf{u}$ reveals that $\phi_2\phi_1[\mathbf{u}] = \psi_2[\mathbf{u}] = \mathbf{w}$. We will consider the set of vertices for which the labels at states $\mathbf{u}$ and $\mathbf{w}$ are different; let $\Delta_{uw} = \{v \in V(\mathcal{G}) \mid \mathbf{u}(v) \neq \mathbf{w}(v)\}$. Since $\psi_2(\mathbf{u}) = \mathbf{w}$, $\Delta_{uw} \subset \text{TR}(\psi_2)$, because the trace is the set of vertex labels that change upon applying ψ_2. However, because $\phi_2\phi_1[\mathbf{u}] = \psi_2[\mathbf{u}]$ and the ϕ_i commute, Δ_{uw} is partitioned by $\text{TR}(\phi_1)$ and $\text{TR}(\phi_2)$ and intersects each set nontrivially. This is because the two generators cannot change any of the same vertex labels (as this would violate commutativity) yet both generators must change some label (as we assumed no generators are trivial). Since $\text{TR}(\phi_1) = \text{TR}(\psi_1)$, we therefore have $\Delta_{uw} \cap \text{TR}(\psi_1) \neq \emptyset$.

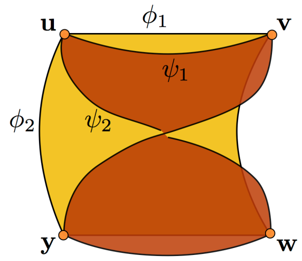

FIGURE 10. Suppose X admits a state structure, with states and generators as shown.

Together these statements imply that $\text{TR}(\psi_1) \cap \text{TR}(\psi_2) \neq \emptyset$, which contradicts the fact that the ψ_i commute. Thus, X cannot be a state complex.

(Y is not state): By observation it is clear that Y is A-special: hyperplanes are simply midplanes and their interactions avoid all the pathologies in Definition 3.6. Suppose, toward a contradiction, that Y admits a state structure. Edges in Y must therefore correspond to generators, so we identify the edges of the base square with generators ϕ_1 and ϕ_2 (Figure 11) and edges of the vertical square with ψ_1 and ψ_2.

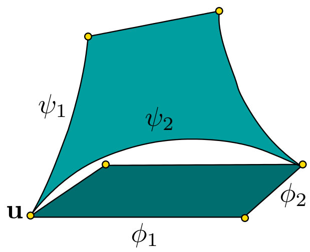

FIGURE 11. Suppose Y admits a state structure, with states and generators as shown.

The presence of 2-cells in Y indicates that ϕ_1 and ϕ_2 commute, as do ψ_1 and ψ_2. In particular, this means

$$\text{SUP}(\phi_1) \cap \text{TR}(\phi_2) = \emptyset \quad \text{and} \quad \text{SUP}(\psi_1) \cap \text{TR}(\psi_2) = \emptyset.$$

The fact that $\psi_2[\mathbf{u}] = \phi_2\phi_1[\mathbf{u}]$ implies that $\text{TR}(\psi_2) = \text{TR}(\phi_2\phi_1) = \text{TR}(\phi_2) \cup \text{TR}(\phi_1)$, by definition of the trace. Since $\text{TR}(\psi_2) \subseteq \text{SUP}(\psi_2)$ we have

$$\text{SUP}(\psi_1) \cap \big(\text{TR}(\phi_2) \cup \text{TR}(\phi_1)\big) = \emptyset,$$

which implies that

$$\text{SUP}(\psi_1) \cap \text{TR}(\phi_2) = \emptyset \quad \text{and} \quad \text{SUP}(\psi_1) \cap \text{TR}(\phi_1) = \emptyset.$$

Thus, ϕ_1 and ϕ_2 change none of the labels on any vertices contained in $\text{SUP}(\psi_1)$. This means that at every state $\mathbf{v}$ where ϕ_1 or ϕ_2 is admissible we have

$$\mathbf{v}|_{\text{SUP}(\psi_1)} = \mathbf{v}^{loc}_{\psi_1} = \phi_i[\mathbf{v}]|_{\text{SUP}(\psi_1)}, \text{ for } i = 1, 2$$

implying that ψ_1 must also be admissible at the states $\phi_1[\mathbf{v}]$ and $\phi_2[\mathbf{v}]$. Specifically, since ψ_1 is admissible at the particular state $\mathbf{u}$ (shown in Figure 11), ψ_1 must apply at states $\phi_1[\mathbf{u}]$ and $\phi_2[\mathbf{u}]$. The complex, however, has no edges at these states representing the action of generator ψ_1, giving a contradiction. Thus, Y cannot be state.

Before proceeding with the last complex, Z, we present a definition and technical lemma that ensure the alphabet for any reconfigurable system can be reduced to a very simple one.

Definition 4.3. *(Isomorphic systems)* We say that two reconfigurable systems are *isomorphic* if there exists a combinatorial isomorphism between their associated state complexes.

Lemma 4.4. *Any reconfigurable system is isomorphic to a reconfigurable system with* $\mathcal{A} = \{0, 1\}$.

Proof. Let $\mathcal{R}$ be a reconfigurable system with domain graph $\mathcal{G}$, alphabet $\mathcal{A}$, generators $\{\phi_i\}$, states $\{\mathbf{u}_t\}$, and state complex $\mathcal{S}$. We may assume without loss of generality that $\mathcal{A} = \mathbb{Z}_n$. We will define a new reconfigurable system, $\bar{\mathcal{R}}$ that is isomorphic to $\mathcal{R}$ and has alphabet $\mathbb{Z}_2$.

We obtain the domain $\bar{\mathcal{G}}$ for $\bar{\mathcal{R}}$ by replacing each vertex $a_i \in \mathcal{G}$ with a copy of K_n, the complete, connected graph on n vertices. We do so by identifying a_i with any vertex $v \in K_n$. The copy of K_n at vertex a_i will be denoted $K_n(i)$ and its vertices denoted $v_0(i), v_1(i), \ldots, v_{n-1}(i)$; these vertex assignments may be made randomly. The vertex $v_j(i) \in \bar{\mathcal{G}}$ will be used to record labelings of the original vertex a_i by element j in $\mathcal{A}$; in this way we pass information from the old alphabet $\mathcal{A}$ to the new graph $\bar{\mathcal{G}}$.

The states are relabeled as follows. Let $\mathbf{u} : V(\mathcal{G}) \to \mathcal{A}$ be a state in $\mathcal{S}$. Each vertex $a_i \in \mathcal{G}$ is assigned some label $l \in \mathcal{A}$ under $\mathbf{u}$. The corresponding state $\bar{\mathbf{u}} : \bar{\mathcal{G}} \to \{0, 1\}$ in $\bar{\mathcal{R}}$ will assign labels to all vertices in $K_n(i)$: $\bar{\mathbf{u}}$ assigns a '1' to vertex $v_l(i)$ and '0' to all others. We redefine generators as follows. For each $\phi \in \mathcal{R}$ we define $\bar{\phi}$ satisfying

$$\begin{aligned} \text{TR}(\bar{\phi}) &= \text{TR}(\phi) \cup \coprod K_n(i) \quad \forall i \text{ such that } a_i \subseteq \text{TR}(\phi), \text{ and} \\ \text{SUP}(\bar{\phi}) &= \text{SUP}(\phi) \cup \coprod K_n(i) \quad \forall i \text{ such that } a_i \subseteq \text{SUP}(\phi). \end{aligned}$$

That is, we add $K_n(i)$ to the trace (respectively, support) of $\bar{\phi}$ whenever a_i is in the trace (support) of ϕ. The local states for $\bar{\phi}$ are the local states for ϕ having been relabeled as per the convention above.

From the definition we see that two generators $\bar{\phi}$ and $\bar{\psi}$ commute if and only if

$$\Big(\textsc{sup}(\phi) \cup \coprod_i K_n(i)\Big) \cap \Big(\textsc{tr}(\psi) \cup \coprod_j K_n(j)\Big) = \emptyset,$$

where $a_i \subseteq \textsc{sup}(\phi)$ and $a_j \subseteq \textsc{tr}(\psi)$. This holds if and only if ϕ and ψ themselves commute. A generator $\bar{\phi}$ is admissible anywhere its local states appear, which is precisely at the relabeled versions of states where ϕ is admissible. It is clear that the two systems produce isomorphic sets of states; together with the previous commutativity and admissibility statements this implies that the corresponding state complexes are isomorphic. □

Remark 4.5. In expanding the domain $\mathcal{G}$ to $\bar{\mathcal{G}}$ the choice of K_n is somewhat arbitrary, as the connectivity of K_n is never used. The proof proceeds in exactly the same manner if we associate n disjoint vertices to $\mathcal{G}$ in place of each original vertex. The graph K_n was chosen for the sake of compact notation and ease of visualization. Indeed, in all abstract examples of reconfigurable systems (i.e., those not representing some physical situation) the edges of $\mathcal{G}$ are inconsequential, as all relabelling actions take place on the vertices.

We now return to the final complex in Theorem 4.1.

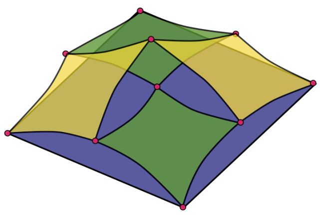

FIGURE 12. Despite uniformity of vertex links (unlike Y) and no twisting (unlike X), the complex Z cannot be endowed with a state structure.

(Z is not state): Suppose, toward contradiction, that Z admits a state structure. We proceed by ruling out, case by case, all possible collections of generators for the underlying reconfigurable system based on the geometry of the complex.

By Lemma 4.4 we may assume that the alphabet is $\mathbb{Z}_2$ and we let $\mathcal{G}$ denote underlying graph associated to this alphabet. We make no general assumptions about $\mathcal{G}$ other than it contains as many vertices as necessitated by the generators we define along the way. By Lemma 5.4 of [20], each hyperplane $\mathcal{H}$ of Z is associated to a unique generator, $\phi_{\mathcal{H}}$. Since $\phi_{\mathcal{H}}$ is nontrivial, it must change the label on (at least) one vertex in $\mathcal{G}$ when it crosses $\mathcal{H}$ in a square containing $\mathcal{H}$ (noting that midplanes and hyperplanes in Z are actually equivalent). No single hyperplane $\mathcal{H}$ of Z separates the space, however. Thus, it is not possible for exactly one generator $\phi_{\mathcal{H}}$ to change the label on a given vertex a in $\mathcal{G}$, else Z would end up with inconsistently labeled states. Hence, if $a \in \text{TR}(\phi)$ for some generator ϕ, then $a \in \text{TR}(\psi)$ for some generator $\psi \neq \phi$ as well.

Next, we assume that $a \in V(\mathcal{G})$ belongs to $\text{TR}(\phi) \cap \text{TR}(\psi)$ and show that there must exist a third generator γ such that a belongs to $\text{TR}(\gamma)$. Ultimately, this will also lead us to inconsistent state labels. Owing to the symmetry of Z and the fact that two generators with overlapping trace must belong to different squares (since they cannot commute), there are only five *distinct* ways to choose generators ϕ and ψ that will change the labels of a vertex a. These cases are illustrated in Figure 13 on a copy of Z that has been cut open at states $\mathbf{u}$ and $\mathbf{v}$ and laid flat.

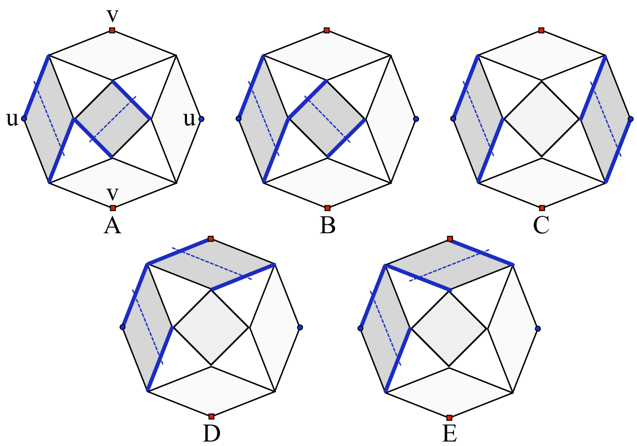

FIGURE 13. In Z, there are five distinct ways to pick generators ϕ and ψ if they have a shared trace element.

Consider case A. Beginning at the leftmost vertex of Z (representing state $\mathbf{u}$), we randomly assign the vertex label '0' to the vertex a and proceed to the right, relabeling a as dictated by the generators ϕ and ψ

(Figure 14, left). If we suppose that no other generators in the complex contain a in their trace, we are forced to continue labeling a as shown in Figure 14 (center). Notice in particular the two circled labels; since a is labeled with a '0' and a '1' at opposite ends of an edge, this edge must correspond to another generator γ which changes a. This forced generator γ is depicted in red in Figure 14 (right). (That γ must exist can also be seen by the fact that, taken together, generators ϕ and ψ as chosen here also fail to separate the space.) Because γ is admissible at vertices where ϕ and ψ are each admissible yet leads to different states, we see that γ is distinct from ϕ and ψ.

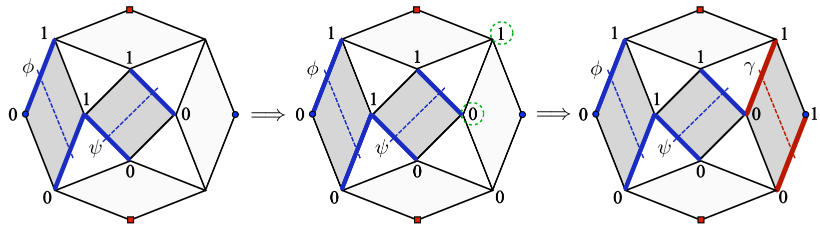

FIGURE 14. Case A, in steps.

Cases B through F are argued similarly. Thus, whenever two generators change the labels on some shared vertex, there must be a distinct third generator that affects that vertex. We will now show that whenever three generators change a shared vertex, the resulting complex Z cannot have consistently labeled states. (Notice that this was true in Case A: the leftmost and rightmost vertices in the complex are actually identified as state $\mathbf{u}$, and therefore must have the same labels for vertex a.)

Figure 15 shows the 8 possible ways (up to symmetry) that three generators, each changing the labels on some vertex a, can be chosen. Note that we do not include the case where the three generators form a closed loop of three edges, as it can be immediately ruled out: each generator must switch the vertex label for a modulo 2, but it is impossible to apply such a switch three times and return to the same labeling.

Consider case $\bar{A}$, the result of investigating case A above. As we noted at the end of case A, when a third generator γ changing the labels on $a \in V(\mathcal{G})$ is present, the resulting complex has labels on the two vertices representing state $\mathbf{u}$ that do not match, at least on a. The presence of other generators in the complex cannot remedy this. This choice of three generators, therefore, is not possible in any state structure on Z.

The same is true of cases $\bar{B}$ through $\bar{F}$; the state $\mathbf{u}$ is labeled inconsistently when we choose three generators in any of these ways.

In the case of $\bar{G}$ we arrive at different contradiction; we label the states of Z as implied by the presence of generators in Figure 16. To achieve the labeling that exists on the center square of the complex, we must have another generator present in the center square that also changes labels on a. This, however, is impossible, because two generators with a shared trace element cannot commute as the center square indicates they do.

Finally, we consider $\bar{H}$. Following the same process, we do not encounter the same contradictions as in earlier cases: the labeling on states is legal (with respect to vertex a, at least) and no pairs of generators are

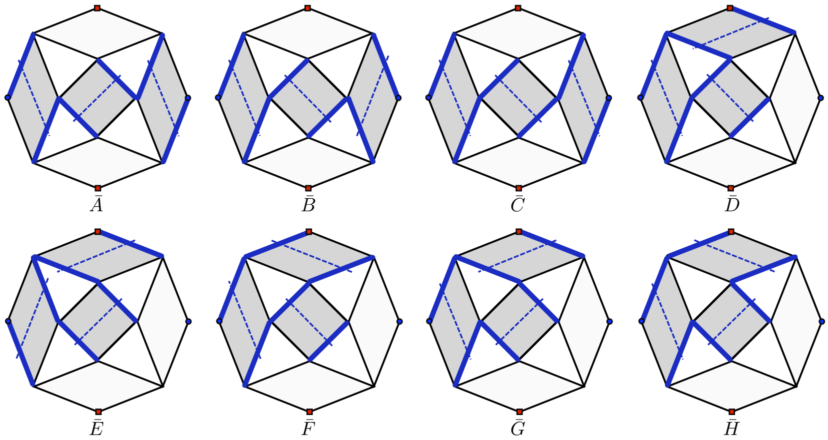

FIGURE 15. Up to symmetry, there are eight ways to select three generators, each of which changes the labels on a fixed vertex a.

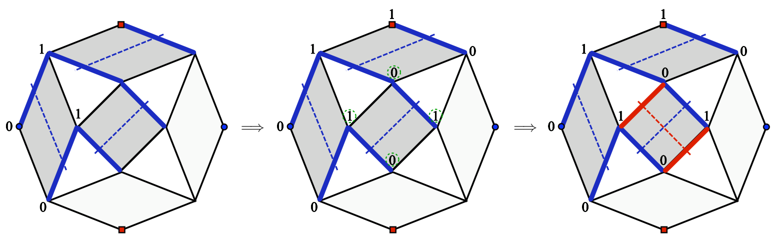

FIGURE 16. Case $\bar{G}$ violates commutativity.

forced to violate commutativity. We may assume that these three generators (i.e., those highlighted in Figure 17, left) are the only generators present that are allowed to change the labels on vertex a. Otherwise, we would have a subcomplex of Z whose generators fall into one of the cases $\bar{A}$ through $\bar{G}$. We therefore fill in the remaining labels on a as forced by these three generators (also Figure 17, left). Now, as state $\mathbf{v}$ (the top and bottom vertices) is distinct from state $\mathbf{u}$, these states must have different labels on some vertex $b \in V(\mathcal{G})$, so without loss of generality we add these labels to Z. Choices for b are written in red in the second component of the label at states $\mathbf{u}$ and $\mathbf{v}$; the empty second slots in Figure 17 indicate that no choice for b has yet been forced.

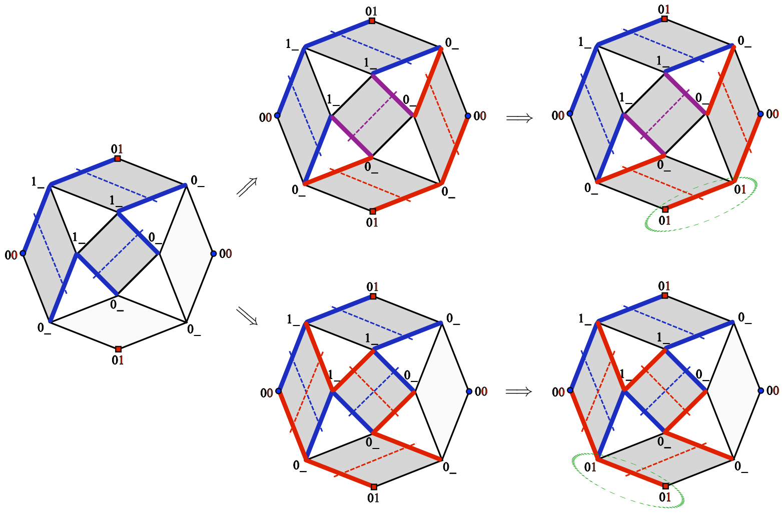

FIGURE 17. Case $\bar{H}$ breaks down into further cases, each of which result in inconsistent labels on Z.

In order to avoid falling into one of the illegal cases above, there is only one allowable arrangement of three generators if each changes the labels on vertex b: that of $\bar{H}$. We may impose this arrangement on the already chosen generators from Case $\bar{H}$ in two ways, shown in Figure 17 (center). In the upper figure, the generator and hyperplane shown in purple (in the central square) represent the fact that there is now a generator in the middle square of the complex that changes labels on both vertices a and b. We fill in the implied labels for vertex b based on these new generators, and in Figure 17 (right) circle the contradiction that arises in labels: two states on opposite ends of a generator changing b cannot have the same label for b.

In summary: No single generator may relabel any vertex of $\mathcal{G}$, and if any two generators relabel a, then there exist at least three generators relabeling a. Having three generators relabel the same vertex, however, results in a set of states inconsistent with the cubical structure on Z. Thus, Z cannot be a state complex, completing the proof of Theorem 4.1. □

Remark 4.6. Discovering that the complex Z was not a state complex was somewhat surprising, as the complex was designed to avoid some of the particular behaviors that had been observed to obstruct 'statehood' in other complexes. For example, there is no twisting, as in the complex X; the link structure is homogeneous, as opposed to the complex Y; and the hyperplanes in the complex do not interact with each other except in single squares, whereas pairs of hyperplanes intersecting in multiple non-adjacent squares within a cube complex had prevented realization as a state complex in a number of other examples considered.

With regards to resolving obstructions to being state, it is notable that all of the non-state examples described above are *virtually* state. That is, all three complexes have finite covers that are state (double covers, in fact). It is also the case that one may make various other alterations to each complex to realize it as a state complex (e.g., adding edges or a square, identifying two states into one, identifying edges), but as the obstructions to being state differ in each example above, so do the corrective changes. We will return to finite covers of state complexes in the next section.

4.2. Operations on state and special cube complexes.

In [22], Haglund and Wise use the notion of walls to prove the following results concerning A-special complexes.

Proposition 4.7 ([22])**.** *Let X and $\{X_i\}$ be A-special cube complexes.*
(1) *Any arbitrary product $\prod_i X_i$ of A-special cube complexes is A-special.*
(2) *Any locally convex subcomplex of an A-special cube complex X is A-special.*
(3) *Any covering space $\hat{X} \to X$ of an A-special cube complex is A-special.*

Here, a subcomplex $X \subset Y$ is *locally convex* if the natural embedding $X \to Y$ is a local isometry. Note that since A-special implies special, all of these results hold when "A-special" is replaced by "special."

Proofs of the above statements are relatively straightforward verifications that pathologies in walls of the new spaces would imply pathologies existed in the walls of the original spaces; see [22] for details.

Property (3) results from the fact that combinatorial maps between cube complexes preserve the special behavior of hyperplanes. Property (2) can be relaxed to include all subcomplexes if one ignores inter-osculating hyperplanes (i.e., *every* subcomplex of an A-special cube complex also avoids self-intersecting, one-sided and self-osculating hyperplanes).

The "locally convex" requirement in Property (2) is a technical condition that is satisfied automatically for the state complexes. For NPC cube complexes X and Y, local convexity can be reduced to the following sufficient condition: for two vertices a and b in $\mathcal{L}k[v] \in X$, if $\phi(a)$ and $\phi(b)$ are adjacent in $\mathcal{L}k[\phi(v)]$, then a and b were adjacent in $\mathcal{L}k[v]$ to begin with [22].

An important step in the classification of state complexes relative to special complexes is to discern in which ways state complexes behave similarly to special complexes (or not). Hence, we now examine analogous statements for state complexes.

4.2.1. **Finite products**. As is the case for special cube complexes, the property "state" is preserved under finite products. When passing to infinite products, although the presence of countably many factors in itself does not prevent the product from being state, we lose some of the underlying structure that was exploited earlier (e.g., the finite shapes condition and a complete, geodesic metric) [9].

Theorem 4.8. *A finite product of state complexes is a state complex.*

Proof. By induction it suffices to show that the product of two state complexes is state. Suppose X and Y are state complexes for reconfigurable systems with, respectively, domain graphs $\mathcal{G}_X$ and $\mathcal{G}_Y$, alphabets $\mathcal{A}_X$ and $\mathcal{A}_Y$, generators $\{\phi_i\}_{i\in\mathcal{I}}$ and $\{\psi_j\}_{j\in\mathcal{J}}$, and states $\{\mathbf{x}_m\}_{m\in M}$ and $\{\mathbf{y}_n\}_{n\in N}$. We must define the reconfigurable system whose state complex $\mathcal{S}$ is isomorphic to $X \times Y$. Informally speaking, to define this system we will essentially take the disjoint union of the underlying systems (though adapting the generators of the individual systems appropriately requires a bit of care).

The alphabet for the system is $\mathcal{A}_X \cup \mathcal{A}_Y$ and the domain graph is the disjoint union $\mathcal{G}_X \amalg \mathcal{G}_Y$. States are ordered pairs of the form $(\mathbf{x}_m, \mathbf{y}_n)$. It is clear that vertices in $X \times Y$ correspond exactly to these states. Generators for the system are of the form $\overline{\phi_i} = (\phi_i, \mathtt{1})$ and $\overline{\psi_j} = (\mathtt{1}, \psi_j)$, where $\mathtt{1}$ denotes the identity in its factor. The supports and traces for generators are defined as follows:

$$\overline{\phi_i} : \begin{cases} \textsc{tr}(\overline{\phi_i}) = \textsc{tr}(\phi_i) \subseteq \mathcal{G}_X \\ \textsc{sup}(\overline{\phi_i}) = \textsc{sup}(\phi_i) \subseteq \mathcal{G}_X \end{cases} , \quad \overline{\psi_j} : \begin{cases} \textsc{tr}(\overline{\psi_j}) = \textsc{tr}(\psi_j) \subseteq \mathcal{G}_Y \\ \textsc{sup}(\overline{\psi_j}) = \textsc{sup}(\psi_j) \subseteq \mathcal{G}_Y . \end{cases}$$

This implies that the generator $\overline{\phi_i}$ is admissible at the states $(\mathbf{x}_m, -)$, where $\mathbf{x}_m$ is some state at which ϕ_i is admissible. Similarly, $\overline{\psi_j}$ is admissible at $(-, \mathbf{y}_n)$, where ψ_j is admissible at $\mathbf{y}_n$. Hence, the action of $\overline{\phi_i}$ on state $(\mathbf{x}_m, \mathbf{y}_n)$ is defined by

$$\overline{\phi_i}\big[(\mathbf{x}_m, \mathbf{y}_n)\big] = (\phi_i, \mathtt{1})\big[(\mathbf{x}_m, \mathbf{y}_n)\big] = \big(\phi_i[\mathbf{x}_m], \mathbf{y}_n\big)$$

which is a state in $X \times Y$ because $\phi_i[\mathbf{x}_m] = \mathbf{x}_{m'}$ is a state in X. Analogously, the action of $\overline{\psi_j}$ on $(\mathbf{x}_m, \mathbf{y}_n)$ is defined by

$$\overline{\psi_j}\big[(\mathbf{x}_m, \mathbf{y}_n)\big] = (\mathtt{1}, \psi_j)\big[(\mathbf{x}_m, \mathbf{y}_n)\big] = \big(\mathbf{x}_m, \psi_j[\mathbf{y}_n]\big).$$

Since an edge in $X \times Y$ is the product of an edge in X with a vertex in Y (or vice versa), every edge in $X \times Y$ corresponds to the action of a generator of the form $\overline{\phi_i}$ or $\overline{\psi_j}$ at an admissible state. Further, since admissibility and generator actions in the system are defined precisely as they were in X and Y, two states $(\mathbf{x}_m, \mathbf{y}_n)$ and $(\mathbf{x}_{m'}, \mathbf{y}_{n'})$ are adjacent in $\mathcal{S}$ if and only if either $(\mathbf{x}_{m'}, \mathbf{y}_{n'}) = (\mathbf{x}_m, \psi_j[\mathbf{y}_n])$ or $(\mathbf{x}_{m'}, \mathbf{y}_{n'}) = (\phi_i[\mathbf{x}_m], \mathbf{y}_n)$. That is, if and only if exactly one component of each state was adjacent in X or Y originally. This gives an isomorphism between the 1-skeleton of $X \times Y$ and the 1-skeleton of $\mathcal{S}$.

To see that the isomorphism extends to higher dimensional cubes, note that the domains $\mathcal{G}_X$ and $\mathcal{G}_Y$ are disjoint, hence $\textsc{tr}(\overline{\phi_i}) \cap \textsc{sup}(\overline{\psi_j}) = \emptyset$ (and vice versa) for all i and j. This implies that $\overline{\phi_i}$ commutes with $\overline{\psi_j}$ whenever they are both admissible; this is precisely at the states $(\mathbf{x}_m, \mathbf{y}_n)$ which admit edges corresponding to the generators ϕ_i and ψ_j. We also see from the definitions that $\overline{\phi_i}$ commutes with $\overline{\phi_{i'}}$ if and only if ϕ_i and $\phi_{i'}$ commute, with a similar statement holding for generators $\overline{\psi_j}$ and $\overline{\psi_{j'}}$. Combining these facts (and reindexing generators if necessary), we arrive at the following:

$$\prod_{i=1}^{l} \phi_i \times \prod_{j=1}^{k-l} \psi_j \text{ is a } k\text{-cube in } X \times Y \iff \prod_{i=1}^{l} \overline{\phi_i} \times \prod_{j=1}^{k-l} \overline{\psi_j} \text{ is a } k\text{-cube in } \mathcal{S},$$

where $l \in \{0, 1, \ldots, k\}$. This establishes the isomorphism between $X \times Y$ and $\mathcal{S}$. □

Remark 4.9. A similar result holds when "finite product" is replaced by "finite wedge product." In fact, the underlying reconfigurable system is even easier to describe for a wedge of two state complexes than for the product above: it is effectively a disjoint union of the two systems, with appropriate adjustments made to ensure commutativity is maintained exactly as in the original complexes. Details appear in [28].

We now demonstrate that, in terms of the properties in Proposition 4.7, the similarities between state and special complexes end here.

4.2.2. **Subcomplexes**. Unlike the property "special," which is preserved when passing to (locally convex) subcomplexes, it is possible to remove cells in a state complex in such a way as to destroy the interactions between generators, resulting in subcomplexes that are not themselves state.

Proposition 4.10. *There exist NPC, locally convex subcomplexes of state complexes that are not state.*

Proof. Recall the non-state complex Y of Theorem 4.1, and let $\overline{Y}$ denote the extended version of Y pictured in Figure 18. Readers may verify that $\overline{Y}$ is a state complex. (One possible reconfigurable system generating $\overline{Y}$ has an underlying graph with three vertices, and there are five distinct generators for the system. Note that both pairs of vertical edges in the complex $\overline{Y}$ must represent a single generator to resolve the obstruction described in the proof of Theorem 4.1.) Checking vertex links confirms that Y is NPC. We verify that Y is a locally convex subcomplex of $\overline{Y}$ by looking at links of vertices of Y under the inclusion $Y \to \overline{Y}$. The only vertices of Y that have different links in $\overline{Y}$ are the two vertices at which the middle square and bottom square of $\overline{Y}$ meet. The link of either vertex in Y is a single edge; in $\overline{Y}$ is two disjoint edges. Since the former is a full subcomplex of the latter, the inclusion $Y \to \overline{Y}$ is a local isometry. Thus, we have found an NPC, locally convex subcomplex Y of a state complex $\overline{Y}$ that is not state. □

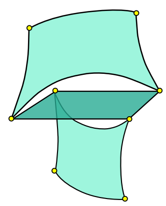

Figure 18. $\overline{Y}$ is a state complex with a subcomplex that does not inherit its state structure.

This example is not pathological; one may construct a number of other examples by deleting subcomplexes of existing state complexes. State complexes do, however, possess a number of natural subcomplexes that are themselves state, including hyperplanes themselves and carriers of hyperplanes [28].

4.2.3. **Finite covers**. We now show that, unlike special complexes, a state complex may possess finite covers that are not state. This is somewhat surprising in light of the fact that all previously examined obstructions to realization can be resolved by taking a two-sheeted cover of the illegal complex. That is, all previous non-state complexes were at least *virtually* state. To see this, we return to our first example of an A-special complex that is not state: the complex X of twisted squares from Theorem 4.1.

It is important to note here that altering the combinatorial structure of X (e.g., subdividing the edges and faces to regard a single square as comprised of four smaller squares, *etc.*) does not resolve the issue of commutativity; conflicts between generators simply propagate to these subdivided generators. In addition, notice that no state complex may contain a copy of X (or any subdivision thereof) as a subcomplex. The presence of this complex in the state complex would imply that the particular generators and states giving rise to it are also present and related in this manner, which we have seen violates the definition of a reconfigurable system. Moreover, the conflict between the supports and traces of two particular generators is independent of their commutativity with other generators, meaning that this obstruction to being state persists even when the squares of X are faces of higher dimensional cubes. Put simply, this relationship between generators cannot exist in any reconfigurable system. We therefore have the following.

Lemma 4.11. *No state complex may contain the complex X of Theorem 4.1 as a subcomplex.*

This fact may be combined with the following "canonical completion and retraction" construction of Haglund and Wise in order to show that state complexes may possess non-state covers.

Theorem 4.12 ([22])**.** *Let B be an NPC special cube complex with simplicial 1-skeleton and let X be any cube complex. If there exists a local isometry $f : X \to B$ then there is a covering $p : C \to B$, an embedding $j : X \to C$ and a cellular map $r : C \to X$ so that $f = pj$ and $rj = 1_X$. That is, there exist maps such that the following diagram commutes:*

$$\begin{tikzcd} & C \arrow[d, "p"] \arrow[dl, shift left, "r"] \\ X \arrow[ur, shift left, "j"] \arrow[r, "f"'] & B \end{tikzcd}$$

Theorem 4.13. *There exists a state complex with a finite cover that is not state.*

Proof. Let X denote the complex of twisted squares (Figure 9, left). Let Γ_X denote the graph whose vertices are distinct hyperplanes $\mathcal{H}$ in X and whose edges correspond to pairs of intersecting hyperplanes in X. We construct the Artin complex associated to Γ_X as in Definition 3.12; in this case, it is possible to describe $\mathcal{A}\textsc{rt}(\Gamma_X)$ explicitly. Since X consists of two squares which share no edges, X has two distinct pairs of intersecting hyperplanes. Thus, $\mathcal{A}\textsc{rt}(\Gamma_X)$ contains one vertex, a loop for each of the four hyperplanes of X, and two squares, disjoint except for their shared vertex. $\mathcal{A}\textsc{rt}(\Gamma_X)$ is therefore a wedge of two 2-tori. Since X is special, there is a local isometry from X to $\mathcal{A}\textsc{rt}(\Gamma_X)$, as per Theorem 3.13.

In order to apply the canonical completion and retraction of Theorem 4.12, we now subdivide $\mathcal{A}\textsc{rt}(\Gamma_X)$ to obtain a cube complex with simplicial 1-skeleton.[7] Adding two vertices to each loop suffices; each loop is now a simplicial cycle of length three. Extending this subdivision to squares in the obvious way yields a complex with a refined cubical structure which is topologically identical to the original complex. This subdivision on $\mathcal{A}\textsc{rt}(\Gamma_X)$ naturally pulls back to X to give an associated subdivision of edges and squares there. By a minor abuse, X and $\mathcal{A}\textsc{rt}(\Gamma_X)$ will now refer to these subdivided complexes.

Note that $\mathcal{A}\textsc{rt}(\Gamma_X)$ is non-positively curved. Prior to subdivision, examining the single vertex v reveals that the link of v is a flag complex; $\mathcal{L}k[v]$ consists of two disjoint cycles with four edges each. Upon subdividing, the link of v remains the same, and it is easily verified that the new vertices of $\mathcal{A}\textsc{rt}(\Gamma_X)$ have flag links as well. Therefore, by Theorem 3.4, $\mathcal{A}\textsc{rt}(\Gamma_X)$ is a state complex: it is an NPC subcomplex of the product of four copies of the (simplicial) graph S^1. (This follows as well from Remark 4.9, since each torus is itself state, being a subcomplex of a product of simple graphs.) This implies in particular that $\mathcal{A}\textsc{rt}(\Gamma_X)$ is A-special. Since both X and $\mathcal{A}\textsc{rt}(\Gamma_X)$ have simplicial 1-skeleta, Theorem 4.12 implies that there exists a covering space $C = C(\mathcal{A}\textsc{rt}(\Gamma_X), X) \to \mathcal{A}\textsc{rt}$ into which X embeds as a subcomplex. That is, the following diagram commutes:

[7] It is possible to avoid this subdivision by instead taking a finite cover of $\mathcal{A}\textsc{rt}(\Gamma_X)$ that has simplicial 1-skeleton. This is because $\mathcal{A}\textsc{rt}(\Gamma_X)$ is a compact, A-special cube complex and therefore has residually finite fundamental group. Such a complex has a CAT(0) universal cover, which must be simplicial, but the residual finiteness yields a finite cover through which the universal cover factors in a way that preserves the simplicial structure (see [22] for details). Showing that a subdivided complex remains a state complex, however, is easier than showing that a particular finite cover is state, so we opt for the subdivision route here.

$$\begin{tikzcd} & C \arrow[dl, shift left, "r"] \arrow[d, "p"] \\ X \arrow[ur, shift left, "j"] \arrow[r, "f"'] & \mathcal{A}\mathrm{RT}(\Gamma_X) \end{tikzcd}$$

Further, C is a finite cover of index $\leq k$, where k is the number of vertices of X [22]. This yields a finite cover of a state complex that contains an embedded copy of the illegal subcomplex X and therefore cannot be state. $\square$

Certainly it is the case that many finite covers of state complexes are again state complexes (in appropriately tidy settings, a "doubling" of the underlying system, with adjustments made to avoid adding extra commutativity, is a tractable solution), but the previous theorem suggests some subtlety exists in any program attempting to classify complexes as virtually state. The hope of realizing some or all special complexes as virtually state stems from the following. As mentioned earlier, special cube complexes were constructed in an effort to generalize clean $\mathcal{VH}$-complexes. A result of Wise [32] shows that all such complexes have some finite cover that embeds in a product of two graphs; hence, all clean $\mathcal{VH}$-complexes are virtually state by Theorem 3.4. Generalizing this directly – showing that all special cube complexes have finite covers that embed in some product of n graphs – seems overly optimistic, but there may be ways to generalize Wise's approach, possibly realizing a subset of state or special complexes as having an inherent graph of spaces structure (as clean $\mathcal{VH}$-complexes have a natural graph of graphs structure, as was exploited in [32]). In any case, it is clear there is much still to be investigated.

References

[1] Aaron David Abrams, *Configuration spaces and braid groups of graphs*. Ph.D. thesis. University of California, Berkeley, 2000.

[2] Aaron Abrams and Robert Ghrist, *Finding topology in a factory: Configuration spaces*, Amer. Math. Monthly **109** (2002), no. 2, 140–150.

[3] ———, *State complexes for metamorphic robots*, Internat. J. Robotics Res. **23** (2004), no. 7-8, 809–824.

[4] Ian Agol, Daniel Groves, and Jason Manning, *The virtual Haken conjecture*. Preprint. Available at arXiv:1204.2810v1 [math.GT].

[5] A. D. Aleksandrov, *Ruled surfaces in metric spaces* (Russian), Vestnik Leningrad. Univ. **12** (1957), no. 1, 5–26, 207.

[6] A. Baudisch, *Subgroups of semifree groups*, Acta Math. Acad. Sci. Hungar. **38** (1981), no. 1-4, 19–28.

[7] Matthew O. Blunt, James C. Russell, María del Carmen Giménez-López, Juan P. Garrahan, Xiang Lin, Martin Schröder, Neil R. Champness, and Peter H. Beton, *Random tiling and topological defects in a two-dimensional molecular network*, Science **322** (2008), no. 5904, 1077–1081.

[8] Martin R. Bridson, *Geodesics and curvature in metric simplicial complexes* in Group Theory from a Geometrical Viewpoint (Trieste, 1990). River Edge, NJ: World Sci. Publ., 1991. 373–463.

[9] Martin R. Bridson and André Haefliger, *Metric Spaces of Non-Positive Curvature.* Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], 319. Berlin: Springer-Verlag, 1999.

[10] Daniel R. Cavagnaro, *Projection of a medium*, J. Math. Psych. **52** (2008), no. 1, 55–63.

[11] Ruth Charney, *An introduction to right-angled Artin groups*, Geom. Dedicata **125** (2007), 141–158.

[12] Gregory S. Chirikjian, *Kinematics of a metamorphic robotic system*, Proc. IEEE ICRA (1994), 449–455.

[13] Michael W. Davis, *Groups generated by reflections and aspherical manifolds not covered by Euclidean space*, Ann. of Math. (2) **117** (1983), no. 2, 293–324.

[14] Michael W. Davis and Tadeusz Januszkiewicz, *Right-angled Artin groups are commensurable with right-angled Coxeter groups*, J. Pure Appl. Algebra **153** (2000), no. 3, 229–235.

[15] Carl Droms, *Graph groups, coherence, and three-manifolds*, J. Algebra **106** (1987), no. 2, 484–489.

[16] ———, *Isomorphisms of graph groups*, Proc. Amer. Math. Soc. **100** (1987), no. 3, 407–408.

[17] ———, *Subgroups of graph groups*, J. Algebra **110** (1987), no. 2, 519–522.

[18] Richard B. Fair, A. Khlystov, Vijay Srinivasan, Vamsee K. Pamula, and Kathryn Nicole Weaver, *Integrated chemical/biochemical sample collection, preconcentration, and analysis on a digital microfluidic lab-on-a-chip platform* in Proc. of SPIE, vol. 5591: Lab-on-a-Chip: Platforms, Devices, and Applications. Ed. Linda A. Smith and Daniel Sobek. 2004. 113–124.

[19] Robert Ghrist, *Configuration spaces and braid groups on graphs in robotics* in Knots, Braids, and Mapping Class Groups—Papers Dedicated to Joan S. Birman (New York, 1998). Ed. Jane Gilman, William W. Menasco, and Xiao-Song Lin. AMS/IP Studies in Advanced Mathematics, 24. Providence, RI: Amer. Math. Soc., 2001. 29–40.

[20] R. Ghrist and V. Peterson, *The geometry and topology of reconfiguration*, Adv. in Appl. Math. **38** (2007), no. 3, 302–323.

[21] M. Gromov, *Hyperbolic groups* in Essays in Group Theory. Ed. S. M. Gersten. Mathematical Sciences Research Institute Publications, 8. New York: Springer, 1987. 75–263.

[22] Frédéric Haglund and Daniel T. Wise, *Special cube complexes*, Geom. Funct. Anal. **17** (2008), no. 5, 1551–1620.

[23] ———, *A combination theorem for special cube complexes*, Ann. of Math. (2) **176** (2012), no. 3, 1427–1482.

[24] John Hempel, 3-*Manifolds.* Annals of Mathematics Studies, 86. Princeton, N. J.: Princeton University Press, 1976.

[25] Tim Hsu and Daniel T. Wise, *On linear and residual properties of graph products*, Michigan Math. J. **46** (1999), no. 2, 251–259.

[26] Darren Long and Alan W. Reid, *Surface Subgroups and Subgroup Separability in 3-Manifold Topology*. Publicações Matemáticas do IMPA [IMPA Mathematical Publications]. Rio de Janeiro: Instituto Nacional de Matemática Pura e Aplicada (IMPA), 2005.

[27] G. A. Niblo and L. D. Reeves, *The geometry of cube complexes and the complexity of their fundamental groups*, Topology **37** (1998), no. 3, 621–633.

[28] Valerie J. Peterson, *On state complexes and special cube complexes*. Ph.D. thesis. University of Illinois at Urbana-Champaign, 2009.

[29] Andrej Šali, Eugene Shakhnovich, and Martin Karplus, *Kinetics of protein folding: A lattice model study of the requirements for folding to the native state*, J. Mol. Bio. **235** (1994), no. 5, 1614–1638.

[30] M. Salvetti, *Topology of the complement of real hyperplanes in* $\mathbb{C}^N$, Invent. Math. **88** (1987), no. 3, 603–618.

[31] Daniel T. Wise, *Non-positively curved squared complexes aperiodic tilings and non-residually finite groups*. Ph.D. thesis. Princeton University, 1996.

[32] ———, *Subgroup separability of the figure 8 knot group*, Topology **45** (2006), no. 3, 421–463.

Department of Mathematics; University of Portland; Portland, OR 97203

E-mail address: petersov@up.edu