

STATE COMPLEXES FOR METAMORPHIC ROBOTS*

A. ABRAMS[†] AND R. GHRIST[‡]

Abstract. A **metamorphic** robotic system is an aggregate of homogeneous robot units which can individually and selectively locomote in such a way as to change the global shape of the system. We introduce a mathematical framework for defining and analyzing general metamorphic robots. With this formal structure, combined with ideas from geometric group theory, we define a new type of configuration space for metamorphic robots — the **state complex** — which is especially adapted to parallelization. We present an algorithm for optimizing an input reconfiguration sequence with respect to elapsed time. A universal geometric property of state complexes — **non-positive curvature** — is the key to proving convergence to the globally time-optimal solution obtainable from the initial path.

1. Introduction. In recent years, several groups in the robotics community have been modeling and building **reconfigurable** or, more specifically, **metamorphic robots** (*e.g.*, [7, 8, 15, 17, 19, 23, 27, 28]). Such a system consists of multiple identical robotic cells in an underlying lattice structure which can disconnect/reconnect with adjacent neighbors, and slide, pivot, or otherwise locomote to neighboring lattice points following prescribed rules: see Fig. 1.1. There are as many models for such robots as there are researchers in the sub-field: 2-d and 3-d lattices; hexagonal, square, and dodecahedral cells; pivoting or sliding motion: see, *e.g.*, [8, 15, 16, 17, 18, 6, 20, 27, 28] and the references therein. The common feature of these robots is an aggregate of lattice-based cells having prescribed local transitions from one shape to another.

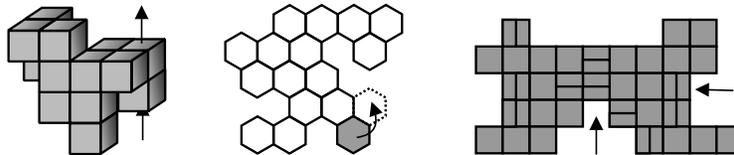


FIG. 1.1. *Metamorphic systems may be built on a variety of lattice structures with sliding or pivoting motion.*

The primary challenge for such systems is **shape-planning**: how to move from one shape to another via legal moves. One centralized approach [9, 10, 22] is to build a **transition graph** whose vertices are the various shapes and whose edges are elementary legal moves from one shape to the next. It is easily demonstrated that the size of this graph is exponential in the number of cells.

This paper extends the notion of a configuration space to metamorphic robots in a novel manner. The idea: consider the transition graph described above as a one-dimensional skeleton of a higher-dimensional cubical complex, the **state complex**. Assume that from a given state there are two legal moves which are physically independent (or, more suggestively, “commutative”): *i.e.*, these moves can be executed simultaneously. In

*AA SUPPORTED BY NATIONAL SCIENCE FOUNDATION GRANT DMS-0089927. RG SUPPORTED BY NATIONAL SCIENCE FOUNDATION GRANT DMS-0134408.

[†] Department of Mathematics, University of Georgia, Athens, GA 30602, USA

[‡] Department of Mathematics, University of Illinois, Urbana, IL 61801, USA

the transition graph, this corresponds to the four edges of a square. For any pair of commutative moves, fill in the four edges of the graph with an abstract square 2-cell. Continue inductively adding k -dimensional cubes corresponding to k -tuples of physically independent motions. The result is a cubical complex which has several advantages over the transition graph:

1. **Simplicity.** The state complex, though larger than the (already large) transition graph, is often topologically much simpler: *e.g.*, the 1-d graph of an n -dimensional cube has $n2^{n-1}$ edges. The number of edges is not necessarily the best measure of complexity: it belies the simplicity of the single cube.
2. **Speed.** Geodesics on this complex cut across the diagonals of cubes whenever possible. One performs all possible commutative motions simultaneously, maximizing parallelization and yielding a speed-up by a factor equal to the number of coordinated motions.
3. **Shape.** The global geometry/topology of the state complex carries information about the metamorphic system. For certain examples, the topology of the state complex “converges” upon refining the lattice. In addition, only special geometries can be realized as the state complex of a local metamorphic system: commutativity in reconfiguration leads to an abhorrence of positive curvature in the state complex.

Sections 2 through 4 give definitions and examples of [abstract] metamorphic systems and their state complexes. The next two sections (Sections 5-6) detail topological and geometric features of the state complex.

For large systems, the problem of computing the state complex and designing geodesics in order to perform shape planning is computationally infeasible, primarily because the size of the complex is often exponential in the number of robot cells. In addition, any control scheme induced by geodesic construction is necessarily centralized. Several researchers have begun building decentralized control algorithms for shape planning [5, 25, 24, 26, 29, 30]. Such algorithms have the advantage of speed and scalability; however, the reconfiguration paths are typically not optimal.

We present an algorithm which works in conjunction with these fast planners to optimize their paths. In Sections 7-8, we present an algorithm for trajectory optimization which takes as its argument an arbitrary edge path in the transition graph. Algorithm 8.2 then performs a type of **curve shortening** within the state complex. A deep theorem about the curvature of all state complexes (Theorem 6.2) is then used to prove that this algorithm returns a shape trajectory which is the global minimum obtainable from this path with respect to elapsed time.

Our definitions and theorems are phrased for systems involving “discrete” reconfiguration. More general types of robots which employ continuous reconfiguration for locomotive gaits (such as the Polybot developed by M. Yim’s lab at PARC) are not covered by our definitions. We note, however, that certain locomotive reconfigurable robots can be thought of as lattice-based tiles by amalgamating subsystems [16]. In addition, our definitions can easily be extended to more general non-lattice reconfigurable systems [3].

2. A mathematical definition. While it is easy to generate examples of what is meant by a metamorphic robot, it is more challenging to write a clean mathematical definition. We propose a set of definitions which is broad enough to include some non-

obvious examples. The paper [6] suggests a similar type of structure using cellular automata rule sets.

A local metamorphic system is a collection of states on a lattice (or, more generally, a graph), where each state is thought of as a labeling function for the aggregate. In most examples, the alphabet of labels will be $\{0, 1\}$ and will be used to denote the absence or presence (resp.) of a module at that lattice point. Any state can be modified by local rearrangements, these local changes being coordinated by a catalogue of models realized under the actions of isometries into the workspace. The adjective “local” refers to legality criteria: anywhere in the workspace at which a local change from the catalogue can be applied, it is legal to do so. To incorporate obstacles and basepoints into our systems, we distinguish between the amount of information needed to determine the legality of an elementary move (the “support” of the move) and the precise place in which modules are actually in motion (the “trace” of the move).

DEFINITION 2.1. *Let \mathcal{L} denote a lattice in \mathbb{R}^k , let $\mathcal{W} \subset \mathcal{L}$ denote a compact workspace, and let \mathcal{A} denote an alphabet of labels. The **catalogue** \mathbf{C} for a local metamorphic system on \mathcal{W} is a collection of **generators**. Each generator $\phi \in \mathbf{C}$ consists of (1) the **support**, $\text{SUP}(\phi) \subset \mathcal{L}$; (2) the **trace** of the move, $\text{TR}(\phi) \subset \text{SUP}(\phi)$; and (3) an unordered pair of local states $\hat{U}_{0,1} : \text{SUP}(\phi) \rightarrow \mathcal{A}$ satisfying¹*

$$(2.1) \quad \hat{U}_0 \Big|_{\text{SUP}(\phi) - \text{TR}(\phi)} = \hat{U}_1 \Big|_{\text{SUP}(\phi) - \text{TR}(\phi)}.$$

Otherwise said, the local states are equal on $\text{SUP}(\phi) - \text{TR}(\phi)$.

DEFINITION 2.2. *An **action** of a generator $\phi \in \mathbf{C}$ is a rigid translation $\Phi : \text{SUP}(\phi) \hookrightarrow \mathcal{W}$. Given a state $U : \mathcal{W} \rightarrow \mathcal{A}$, an action Φ is said to be **admissible** at U if $\hat{U}_0 = U \circ \Phi$. In this case, we write*

$$\phi[U] := \begin{cases} U & : \text{ on } \mathcal{W} - \Phi(\text{SUP}(\phi)) \\ \hat{U}_1 \circ \Phi^{-1} & : \text{ on } \Phi(\text{SUP}(\phi)) \end{cases}.$$

(Note that Φ is left out of the notation.)

DEFINITION 2.3. *A **local metamorphic system** on \mathcal{W} is a collection of states $\{U_\alpha : \mathcal{W} \rightarrow \mathcal{A}\}$ closed under all possible admissible actions of generators in the catalogue \mathbf{C} .*

To repeat, the catalogue and the workspace are the “seeds” for a local metamorphic system. From this pair, all possible translations of the supports into \mathcal{W} yield the actions. Then, a collection of states on the workspace is a local metamorphic system if, whenever an action Φ of a generator ϕ on a state U is admissible, then the corresponding state $\phi[U]$ is also included.

A metamorphic system with **obstacles** $\mathcal{O} \subset \mathcal{W}$ satisfies in addition

$$(2.2) \quad \Phi(\text{TR}(\phi)) \cap \mathcal{O} = \emptyset,$$

for each $\phi \in \mathbf{C}$. Obstacle sets count as legal positions for determining the admissibility of a move (the support of an action may intersect \mathcal{O}), but no motion of metamorphic agents may incorporate the obstacle sites (the trace of an action must not intersect \mathcal{O}).

¹All generators are assumed to be **nongenerate** in the sense that $\hat{U}_0 \neq \hat{U}_1$.

3. Examples. Some of the following examples are inspired by metamorphic robots already developed; other examples are more abstract. Unless otherwise noted, the alphabet \mathcal{A} is $\{0, 1\}$, denoting unoccupied or occupied cells.

EXAMPLE 3.1. [2-d hex with pivots] We present two slightly different catalogues, each with six generators (or one, up to discrete rotations), in Fig. 3.1. Both of these systems, modeled after that of [8], have local moves which pivot a planar hexagon about a neighbor. For all generators presented, the trace is equal to the two central hexagons. In the first system, the support is chosen so that the aggregate does not change its topology but only its shape. The slightly smaller support of the second catalogue allows for local topology changes. To model a fixed “base” cell (which is, say, affixed to a power source as in [20]), one establishes this cell as an obstacle \mathcal{O} .

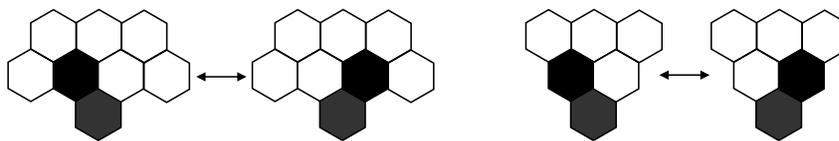


FIG. 3.1. Two different catalogues for a 2-d hexagonal lattice system with pivots. Black cells are occupied, white are unoccupied. The local states U_0 and U_1 are shown for each generator.

EXAMPLE 3.2. [2-d square lattice] In Fig. 3.2, we display a generator for a planar system in which rows [as pictured] and columns [not pictured] of an aggregate of square cells can slide. There are in fact several generators represented in “shorthand,” one for each $k \geq 0$. A dot inside a cell indicates that it can be either occupied or unoccupied, but if occupied, then its neighbor (indicated by an arrow) must also be occupied. This condition guarantees that the aggregate does not disconnect (even locally) under slides. The trace of this set of generators is the entire middle row except the two endpoints. To keep the catalogue finite, one would include only those generators with $k \leq N$, where N is the number of occupied cells in any state.

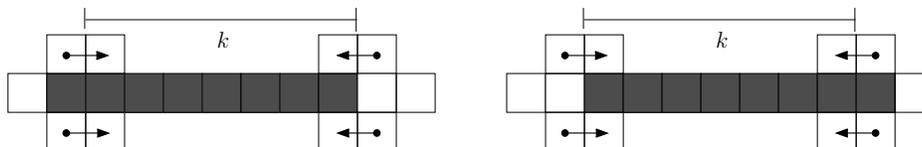


FIG. 3.2. The row-sliding generators for a sliding-squares system.

EXAMPLE 3.3. [2-d articulated planar arm] Consider as a workspace \mathcal{W} the set of edges in the planar integer lattice. The catalogue consists of two generators, pictured in Fig. 3.3. Beginning with a state having N vertical edges end-to-end, the metamorphic system thus generated models the position of an articulated robotic arm with fixed base which can (1) rotate at the top end and (2) flip corners as per the diagram.

If one includes rotations of these generators, more intricate types of configurations are possible, including deadlocked configurations. Examples with multiple interacting arms can be realized by adding new generators: an “attach-detach” generator which allows endpoints of arms to merge (thus yielding a “marked point” at the attachment); and a “sliding” generator which allows this marked point to slide, having the effect of allowing

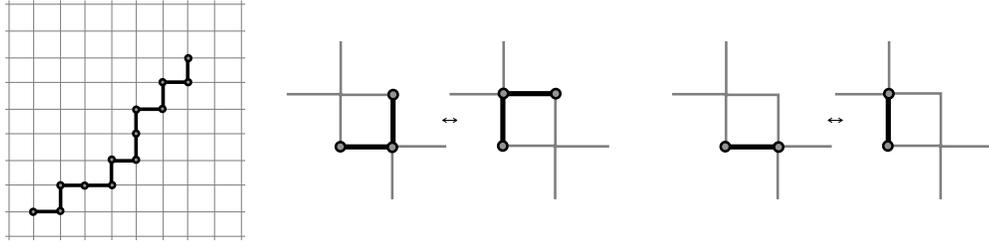


FIG. 3.3. A positive articulated robot arm [left]. One generator [center] flips corners and has as its trace the central four edges. The other generator [right] rotates the end of the arm, and has trace equal to the two activated edges.

the attached arms to trade segments. These generators are pictured (up to Euclidean symmetries) in Fig. 3.4.

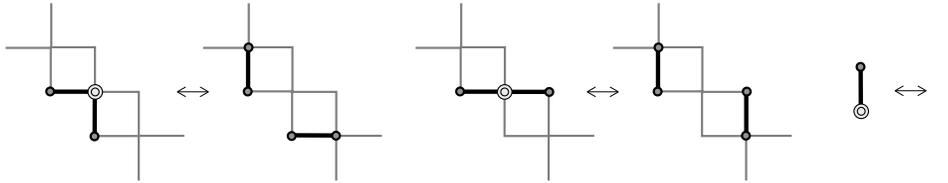


FIG. 3.4. Additional generators allow for attachment and detachment of arm endpoints [left, center] and sliding of the coupling points [right].

EXAMPLE 3.4. [2-d expansion-compression square system] Consider a planar square lattice workspace. This system will use an alphabet of labels $\mathcal{A} = \{0, 1, =, ||\}$ whose interpretation is as follows: “0” means that a cell is unoccupied; “1” means that the cell is occupied by one module; “=” and “||” imply that the cell is occupied by two modules compressed together in a horizontal or vertical orientation (resp.). The catalogue consists of six generators, illustrated in Fig. 3.5 (the lower two generators are only represented up to flips). The trace is equal to the support for the top two generators illustrated; for the bottom two generators, the trace is equal to the support minus the single square which remains unoccupied (label “0”).

This example is a local system based on the Crystalline robots of Rus et al. [23]. Extensions to 3-d cubical systems and more elaborate motions can be accommodated with minor modifications.

One of the benefits of writing down a rigorous definition of a metamorphic system is the discovery of systems which have little resemblance to the systems of, say, Fig. 1.1. In particular, our definitions easily extend to metamorphic systems which are not lattice-based. The following example is especially interesting.

EXAMPLE 3.5. Consider a finite graph Γ in which every edge is assigned a length of one. (Every graph can be embedded in some \mathbb{R}^k so as to have this property.) The catalogue consists of a single generator whose support and trace are precisely the closure of a single abstract edge. The local states of this generator consist of the pair \hat{U}_0 and \hat{U}_1 which evaluate to 1 on one of the endpoints and 0 on the other. The actions in this case are length-preserving maps from the abstract edge into Γ . The metamorphic system

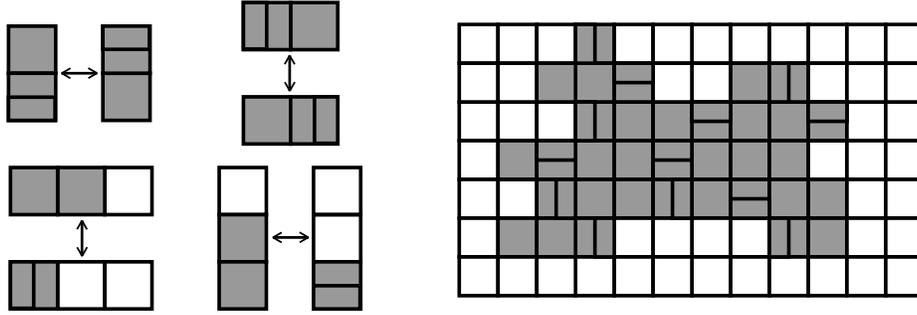


FIG. 3.5. Generators for a simple compression-expansion system on the square planar lattice [left]; an example of a typical state [right].

generated from a state U on Γ with N vertices evaluating to 1 mimics an ensemble of N unlabeled non-colliding Automated Guided Vehicles on Γ , cf. [13]. To work with labeled vehicles, one need merely increase the alphabet \mathcal{A} to accommodate the different labels.

More abstract examples of metamorphic examples include spaces of triangulations of polygons with edge-flipping as the generator, examples arising from word representations in group theory, and certain multi-step assembly processes [3].

4. The state complex. In the robotics literature, one often models a configuration space for a metamorphic system with a transition graph which represents actions of elementary moves on states. That is, the vertex set is the collection of all states $\{U_\alpha\}$, and the edges are unoriented pairs of states which differ by the action of one generator. Transition graphs are discussed for shape-planning in several particular cases in the literature (planar hex case: [9, 22]). Our departure is to make the transition graph the 1-skeleton of a cubical complex (an analogue of a simplicial complex, but made out of abstract cubes) which coordinates parallel or “commutative” motions.

DEFINITION 4.1. *In a local metamorphic system, a collection of actions of (not necessarily distinct) generators $\{(\phi_{\alpha_i}, \Phi_{\alpha_i})\}$ is said to **commute** if*

$$(4.1) \quad \Phi_{\alpha_i}(\text{TR}(\phi_{\alpha_i})) \cap \Phi_{\alpha_j}(\text{SUP}(\phi_{\alpha_j})) = \emptyset \quad \forall i \neq j.$$

EXAMPLE 4.2. Two simple examples suffice to illustrate the difference between commuting and noncommuting actions. First, consider the pair of commuting moves for a planar hexagonal pivoting system, as represented in Fig. 4.1 [left]. Compare this with a planar sliding block example as illustrated in Fig. 4.1 [right]. Although the pair of moves illustrated forms a square² in the transition graph, this particular pair of actions does not commute. Physically, it is obvious why these moves are not independent: sliding the column part-way obstructs sliding a transverse row. Mathematically, this is captured by the traces of the actions intersecting.

The state complex has an abstract k -cube for each collection of k admissible commuting actions:

²The individual robotic cells are not labeled: only the shape of the aggregate is recorded.

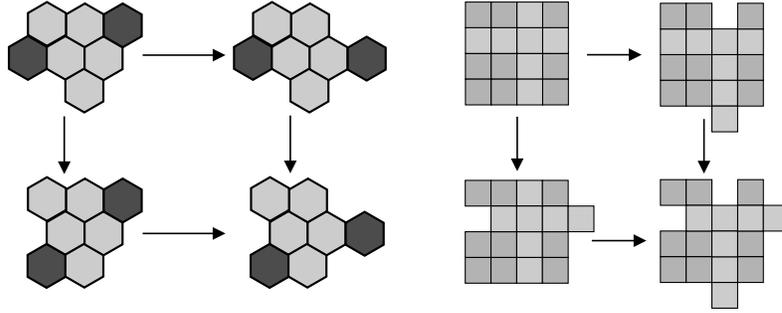


FIG. 4.1. Examples of commuting [left] and noncommuting [right] actions in planar systems.

DEFINITION 4.3. The **state complex** \mathcal{S} of a local metamorphic system is the following abstract cubical complex. Each abstract k -cube $e^{(k)}$ of \mathcal{S} is an equivalence class $[U; (\Phi_{\alpha_i})_{i=1}^k]$ where

1. $(\Phi_{\alpha_i})_{i=1}^k$ is a k -tuple of commuting actions of generators ϕ_{α_i} ;
2. U is some state for which all the actions $(\Phi_{\alpha_i})_{i=1}^k$ are admissible; and
3. $[U_0; (\Phi_{\alpha_i})_{i=1}^k] = [U_1; (\Phi_{\beta_i})_{i=1}^k]$ if and only if the list (β_i) is a permutation of (α_i) and $U_0 = U_1$ on the set $\mathcal{W} - \bigcup_i \text{SUP}(\phi_{\alpha_i})$.

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

$$(4.2) \quad \partial[U; (\Phi_{\alpha_i})_{i=1}^k] = \bigcup_{i=1}^k ([U; (\Phi_{\alpha_j})_{j \neq i}] \cup [\phi_{\alpha_i}[U]; (\Phi_{\alpha_j})_{j \neq i}]).$$

It follows easily that the k -cells are well-defined with respect to admissibility of actions. The proof of the following obvious lemma is given in detail to flesh out the previous definition.

LEMMA 4.4. (a) The 0-dimensional skeleton of \mathcal{S} , $\mathcal{S}^{(0)}$, is the set of states in the reconfigurable system. (b) The 1-dimensional skeleton of \mathcal{S} , $\mathcal{S}^{(1)}$, is precisely the transition graph.

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

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

For small numbers of cells, it is easy to illustrate the state complex.

EXAMPLE 4.5. Consider the 2-d square lattice row/column sliding system whose catalogue is illustrated in Fig. 3.2. If we consider a system with fixed obstacles in the form of

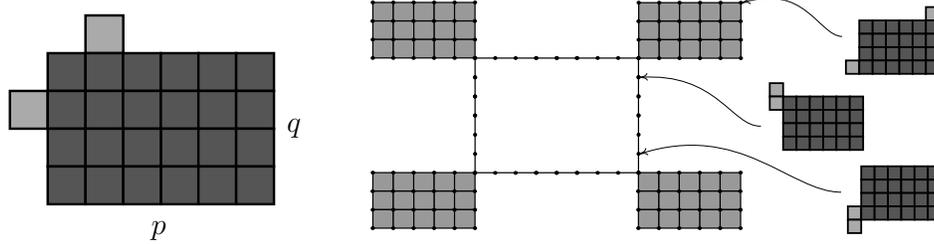


FIG. 4.2. The planar sliding square example [left] with two movable blocks [in grey] and a p -by- q obstacle set [in black] yields a state complex \mathcal{S} that is topologically a circle [right]. Sample configurations in the state complex are illustrated.

a p -by- q rectangle generated from the state of Fig. 4.2 [left], one obtains a planar transition graph with $4(pq + 1) + 2(p + q)$ vertices and $8(pq + 1) - 2(p + q)$ edges. In contrast, the state complex is that of Fig. 4.2 [right]: this is topologically a circle, corresponding to the fact that the pair of free squares can circulate about the obstacle set through a sequence of slides. The large 2-d regions correspond to states in which the two free squares are on separate (but adjacent) sides of the obstacle set.

EXAMPLE 4.6. Consider the planar hex system of Fig. 3.1 [left] with a workspace \mathcal{W} consisting of a long channel of four rows with a line of occupied cells attached to a fixed obstacle at the right. This line of cells can “climb” on itself from the left and migrate to the right, one by one. The entire state complex is illustrated in Fig. 4.3[right]. Although the transition graphs appear complicated, this state complex is contractible for any length channel.

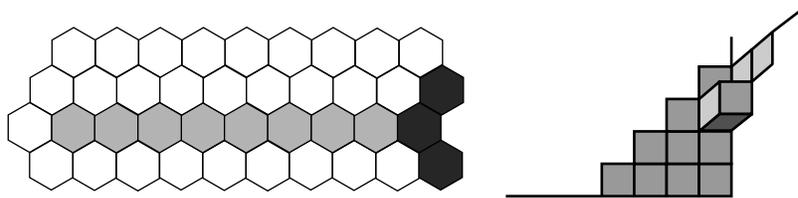


FIG. 4.3. For a line of hexagons filing out of a constrained tunnel [left], the state complex is contractible [right]. Black cells represent obstacle sets. A short tunnel is shown: for longer tunnels, the state complex has higher dimension, but is still contractible.

EXAMPLE 4.7. The state complex associated to the positive articulated robot arm of Example 3.3 in the case $N = 5$ is given in Fig. 4.4. Note that there can be at most three independent motions (when the arm is in a “staircase” configuration); hence the state complex has top dimension three. Notice also that although the transition graph for this system is complicated, the state complex itself is topologically trivial (contractible).

EXAMPLE 4.8. In the system of Example 3.5 with the graph being a K_5 (the complete graph on five vertices) and $N = 2$, the state complex is a two-dimensional closed surface. A simple combinatorial argument (as in [1, 2]) reveals that the Euler characteristic is -5 , implying that the state complex is non-orientable. If the AGV’s are labeled, the state complex becomes a closed orientable surface of genus 6.

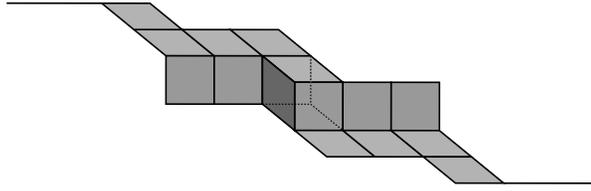


FIG. 4.4. *The state complex of a 5-link positive arm has one cell of dimension three, along with several cells of lower dimension.*

5. The topology of \mathcal{S} . If one looks at a transition graph without knowing the particulars of the metamorphic system, very little information can be extracted. This paper argues that completing the transition graph to the state complex is “natural” — the state complex simplifies the transition graph and endows it with topological and geometric content.

Our first example of naturality is motivated by the desire to build metamorphic systems with large numbers of micro- or nano-scale cells. While large numbers of cells would yield a type of continuum-limit convergence on the dynamics of shape change, the resulting transition graphs have no such convergence. The size of the transition graph goes up exponentially in the number of cells; more ominous is that the topology of the transition graph (the number of basic cycles) blows up as well. This is not always so with the state complex: in certain key examples, the topology of \mathcal{S} is either invariant or converges to a limiting type.

We have already seen one such example of this stabilization. In Example 4.5, the state complex of a pair of squares sliding along a rectangular obstacle of size p -by- q is topologically a circle, independent of p and q . This can be interpreted as a type of convergence: consider the effect of refining the underlying lattice structure, increasing p and q while maintaining a pair of sliding squares. Then the dimension of the state complex remains the same, as does the topological type of the space. Intuitively speaking, the “limit” as this refining process is repeated yields a “topological” configuration space of two points sliding smoothly along the boundary of a rectangle, which can dock or un-dock at the corners.

EXAMPLE 5.1. Recall the state complex associated to the metamorphic system of N points on a graph Γ , Example 3.5. Consider a refinement of Γ which inserts additional vertices along edges. It follows from the techniques of [1] that the state complex of this refined system has the same topological type (up to homotopy equivalence) after a fixed bound on the refinement (N additional vertices per edge). Furthermore, this “stabilized” state complex is in fact homotopic to the topological configuration space of N non-colliding points of Γ — precisely what one expects as the number of refinements goes to infinity.

EXAMPLE 5.2. Recall the positive articulated robot arm of Example 3.3. Consider a refinement of the underlying lattice which shrinks the lattice by a factor of two (or, equivalently, which inserts an additional joint in the middle of each edge). This is a more dramatic change since the dimension of the state complex doubles. Nevertheless, the topological type is invariant: the state complex remains contractible.

PROPOSITION 5.3. *Let \mathcal{S}_N denote the state complex of the positive articulated arm from*

Example 3.3 with N segments. The complexes \mathcal{S}_N are all contractible.

PROOF: For these articulated arms, there is a nice inductive structure on the state complexes. Fixing N , each state (vertex) in \mathcal{S}_N is represented as a length N word in the symbols x and y , where x denotes the arm going to the right and y denotes the arm going up. In this language, the two generators are (1) transposing a subword $xy \leftrightarrow yx$, and (2) changing the last letter of the word.

Consider the subcomplex $X \subset \mathcal{S}_N$ consisting of all cells whose vertices have words beginning with the letter x . Likewise, let Y denote the subcomplex all of whose vertices begin with the letter y . These subcomplexes are each a copy of \mathcal{S}_{N-1} which we may assume inductively is contractible. One passes between the subcomplexes X and Y only when a move exchanges the initial two letters of the word from xy to yx . The connecting set is thus homeomorphic to $\mathcal{S}_{N-2} \times [0, 1]$ and attached to X and Y along $\mathcal{S}_{N-2} \times \{0\}$ and $\mathcal{S}_{N-2} \times \{1\}$ respectively. Again, by induction, these sets are contractible. A pair of contractible sets joined along contractible subsets is contractible. \diamond

This should come as no surprise: in the limit as $N \rightarrow \infty$, the metamorphic system approximates the configuration space of a smooth curve of fixed length which is positive in the sense that the curve is always nondecreasing in the horizontal and vertical components. That the (infinite dimensional) space of such smooth curves is contractible is easily demonstrated: given any such curve with endpoint fixed at the origin in the plane, pull the other end along the straight line connecting it to the origin until the strand is taut. Then, rotate the line segment rigidly until it is, say, vertical. This is a continuous deformation on the space of all smooth positive curves of fixed length to a single vertical segment.

It is certainly not the case that an arbitrary reconfigurable system possesses such convergence properties: the manner in which one refines the states is important. Still, it appears that state complexes can often be viewed as “discretizations” of some underlying smooth configuration space. This is an important focus for future inquiry.

6. The geometry of \mathcal{S} . There are several natural ways to measure distances in state complexes. We first discuss the geometry arising from considering each cube of \mathcal{S} to be Euclidean (i.e., flat), with unit side length; we call \mathcal{S} with this metric a *Euclidean cube complex*. However, this does not imply that the complex, as a whole, is flat. Indeed, non-zero curvature can be concentrated at places where several cells meet. A simple example appears in Fig. 6.1: here, a surface built from flat 2-cells can be seen to have curvature which depends on the number of 2-cells incident to a vertex. Four incident cells implies zero curvature; three cells implies positive curvature; and five or more cells implies negative curvature. For a two-dimensional complex, this is equivalent to computing the total angle about a vertex.

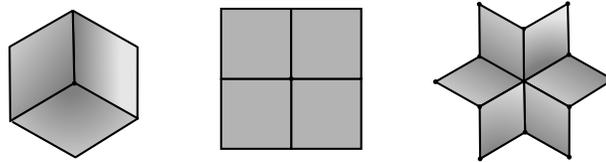


FIG. 6.1. Curvature about a vertex in a 2-d Euclidean cubical complex: positive, zero, and negative.

Such an extension of curvature to general metric spaces is made precise in Gromov's work on curved metric spaces [14] (extending the classical work of Alexandrov, Busemann, and others) in which triangles with geodesic edges are used to measure curvature bounds. In brief, let X be a metric space and $p \in X$ a point. To bound the curvature of X at p , consider a small triangle T about p with geodesic edges of length a , b , and c . Build a **comparison triangle** T' in the Euclidean plane whose sides also have length a , b , and c respectively. Choose a geodesic chord of T and measure its length d . In T' , measure the length d' of the chord whose endpoints correspond to those of the chord in X .

DEFINITION 6.1. A metric space X is **nonpositively curved** (or **NPC**) if for every sufficiently small geodesic triangle T and for every chord of T , it follows that $d \leq d'$.



FIG. 6.2. Comparison triangles measure curvature bounds.

In other words, geodesic chords are no longer than Euclidean comparison chords. It should be stressed that the NPC property is very special and highly desirable. Indeed, being NPC implies a variety of topological consequences reminiscent of smooth nonpositively curved manifolds.

Despite the variety of (local) metamorphic systems, all state complexes share this special geometric property.

THEOREM 6.2. The state complex \mathcal{S} of any local metamorphic system is nonpositively curved.

The proof of this theorem is simple, but requires some additional machinery.

DEFINITION 6.3. Let X denote a complex (either simplicial or cubical) and let v denote a vertex of X . The **link** of v , $\ell k[v]$, is defined to be the abstract complex which has one k -dimensional simplex for each $(k+1)$ -dimensional cube in X incident to v . The boundary relations are those inherited from X : namely, the boundary of a k -simplex in $\ell k[v]$ represented by a $(k+1)$ -cube in X is the set of all simplices represented by the faces of the $(k+1)$ -cube.

Links can be thought of as a simplicial version of the locus of points a small fixed distance from the vertex v .

DEFINITION 6.4. A Euclidean cube complex X satisfies the **link condition** if, for each vertex $v \in X$, $\ell k[v]$ satisfies the following: for each k , if any $k+1$ vertices in $\ell k[v]$ are pairwise connected by edges in $\ell k[v]$, then those vertices bound a unique k -simplex in $\ell k[v]$.

An important and deep theorem of Gromov [14] asserts that a Euclidean cube complex is nonpositively curved if and only if it satisfies the link condition. This criterion makes it easy to prove Theorem 6.2.

PROOF OF THEOREM 6.2: Let U denote a vertex of \mathcal{S} . Consider the link $\ell k[U]$. The 0-cells of the $\ell k[U]$ correspond to all edges in \mathcal{S} incident to U ; that is, actions of generators admissible at the state U . A k -cell of $\ell k[U]$ is thus a commuting set of $k + 1$ of these actions based at U . The interpretation of the link condition for a state complex is as follows: if at U one has a set of $k + 1$ admissible actions, $\{(\phi_{\alpha_i}, \Phi_{\alpha_i})\}_{i=1}^{k+1}$, of which each pair commutes, then the full set of $k + 1$ generators must commute [existence of the k -simplex in the link]. Furthermore, they must commute in a unique manner [uniqueness of the k -simplex in the link].

The proof of existence follows directly from Definition 4.1: any collection of pairwise commutative actions is totally commutative. We therefore have a $(k + 1)$ -dimensional cell in \mathcal{S} which is the equivalence class $[U; (\Phi_{\alpha_i})_{i=1}^{k+1}]$. This is the representative in \mathcal{S} of the k -simplex in $\ell k[U]$. To show uniqueness, consider any other k -simplex in $\ell k[U]$ which has the same vertex set. This must correspond to a $(k + 1)$ -dimensional cube in \mathcal{S} with actions $(\phi_{\alpha_i}, \Phi_{\alpha_i})_{i=1}^{k+1}$ (up to some permutation) based at U . From Definition 4.3, this cell must be the same equivalence class, namely $[U; (\Phi_{\alpha_i})_{i=1}^{k+1}]$. \diamond

EXAMPLE 6.5. To see how the NPC property can fail, consider the following non-local metamorphic system. Recall the generator for the planar hexagonal system presented in Fig. 3.1[right] which (along with its rotations) allows for local disconnection of the aggregate. If we add to this system a global rule that requires the aggregate to be connected (for, say, considerations of power transmission), then we no longer have a local system, and positive curvature may exist. Fig. 6.3 shows a configuration in which three actions of local generators act to disconnect the aggregate locally but not globally. These actions commute pairwise, and any two do not disconnect globally. However, performing all three actions disconnects the space and leads to an illegal state. Therefore, the corresponding state complex for this non-local system has a ‘‘corner’’ of positive curvature as a local factor. Note: the state complex will not be two-dimensional here, but will be locally a product of this with another cubical complex. The positive curvature persists.

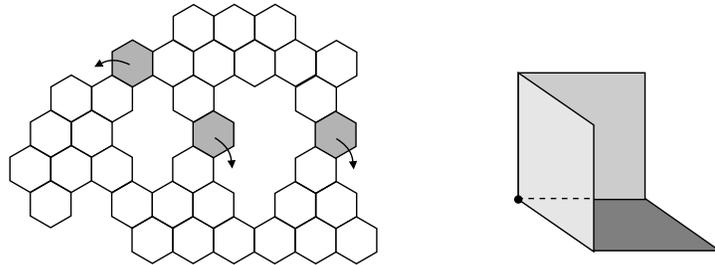


FIG. 6.3. A generator which allows for local disconnection [left] admits configurations [center] for which pairs of actions leave the aggregate connected, but triples do not; this non-local system has positive curvature in the state complex at this vertex [right].

7. Geodesics and time on \mathcal{S} . Besides displaying a unifying geometric feature, the nonpositive curvature has implications for path-planning, and, hence, the shape-planning problem.

COROLLARY 7.1. *Each homotopy class of paths connecting two given points of a state complex contains a unique shortest path.*

PROOF: This is well-known for NPC spaces. The only difficulty lies in proving that the $d \leq d'$ inequality for sufficiently small triangles implies the $d \leq d'$ inequality for *all* geodesic triangles which are contractible in the space: see [4]. Assume, then, that this inequality holds in general, and consider a pair of distinct homotopic shortest paths from points p to q in \mathcal{S} . Choose any point r on one of the two paths. The two halves of this path bisected by r are themselves shortest paths from p to r and r to q respectively. Thus, there is a geodesic triangle in \mathcal{S} whose comparison triangle in the Euclidean plane is a degenerate straight line from p' to q' . The $d \leq d'$ inequality applied to the segment from r to its corresponding point on the other geodesic shows that these points coincide; therefore the two shortest paths from p to q in \mathcal{S} are identical. \diamond

Fig. 7.1[left] gives a simple example of a 2-d cubical complex with positive curvature for which the above corollary fails.

This corollary is a key ingredient in the applications of NPC geometry to path-planning on a configuration space, since one expects geodesics on \mathcal{S} to coincide with optimal solutions to the shape planning problem. However, in the context of robotics applications, the goal of solving the shape-planning problem is *not* necessarily coincident with the geodesic problem on the state complex. Fig. 7.2[left] illustrates the matter concisely. Consider a portion of a state complex \mathcal{S} which is planar and two-dimensional. To get from point p to point q in \mathcal{S} , any edge-path which is weakly monotone increasing in the horizontal and vertical directions is of minimal length in the transition graph. The true geodesic is, of course, the straight line, which is not well-positioned with respect to the discrete cubical structure.

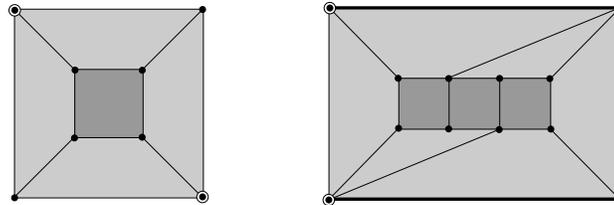


FIG. 7.1. [left] A 2-d cube-complex (with positive curvature) possessing two distinct homotopic shortest paths between a pair of marked points; [right] a 2-d cube-complex (with positive curvature) possessing an edge path (the thick line on the boundary) which is a locally (but not globally) shortest path. Any cube path near this path is strictly longer. Note: for both of these complexes, all edges are assigned unit length.

Given the assumption that *each elementary move can be executed at a uniform maximum rate*, it is clear that the true geodesic on \mathcal{S} is **time-minimal** in the sense that the elapsed time is minimal among all paths from p to q . However, there is an envelope of non-geodesic paths which are yet time-minimizing. Indeed, the true geodesic in Fig. 7.2[left] “slows down” some of the moves unnecessarily in order to maintain the constant slope.

This leads us to define a second metric on \mathcal{S} , one which measures elapsed time. Namely, instead of the Euclidean metric on the cells of \mathcal{S} , consider the space \mathcal{S} with the ℓ^∞ norm on each cell. (This is also called the supremum norm: a vector is measured by the maximum of its components in each coordinate direction.) The geodesics in this geometry represent reconfiguration paths which are *time minimizing*. Using the results of [21], one can prove that these geodesics are easily described using the notion of a **cube**

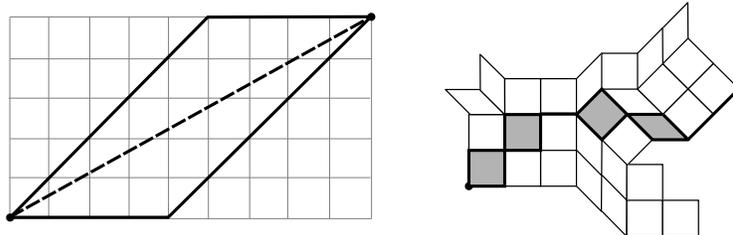


FIG. 7.2. [left] The true geodesic lies within an envelope of time-minimizing paths. No minimal edge-paths are time-minimizing; [right] a normal cube-path in an NPC complex (after [21]).

path.

DEFINITION 7.2. A **cube path** from a vertex v_0 to a vertex v_N is an ordered sequence of closed cubes $\{C_i\}_1^N$ in \mathcal{S} which satisfy (1) $C_i \cap C_{i+1} = v_i$ for some vertex v_i ; and (2) C_i is the smallest cell of \mathcal{S} containing v_{i-1} and v_i . A cube path is said to be **normal** if in addition, (3) for all i , $\text{St}(C_i) \cap C_{i+1} = v_i$, where $\text{St}(C_i)$ is the **star** of C_i (the union of all closed cubes, including C_i , which have C_i as a face).

Roughly speaking, a normal cube path is one which uses the highest dimensional cubes as early as possible in the path.

THEOREM 7.3. In any state complex \mathcal{S} , there exists a unique normal cube path from p to q in each homotopy class. The chain of diagonals through this cube path minimizes time among all paths in its homotopy class.

This result follows from Theorem 6.2 and a result of [21] (or, see the constructive algorithm of Section 8 below for a proof of existence).

This result also implies that, in the category of homotopic cube-paths, there is no such thing as a strictly local minimum for length. See Fig. 7.1[right] for a simple example of a cube complex with some positive curvature that has a cube path which is locally of minimal length among cube paths, but not globally so.

This is very significant. Beginning with any path in the state complex, say, one obtained via a fast distributed algorithm for shape-planning, one can employ a gradient-descent curve shortening on the level of cube paths. The above results imply that any algorithm which monotonically reduces length must converge to the shortest path (in that homotopy class) and cannot be hung up on a locally minimal cube path. The presence of local minima is a persistent problem in optimization schemes: nonpositive curvature is a handy antidote.

8. Optimizing paths. In cases where the state complex is sparse — local principal cells being of high dimension with few neighbors — it is possible to compress the complex significantly. However, shape-planning via constructing all of \mathcal{S} and determining geodesics is, in general, computationally infeasible: the total size of the state complex is often exponential in the number of cells in the aggregate. We therefore assume that some shape trajectory has been determined (perhaps by one of the many excellent decentralized planners available [5, 6, 24, 25, 26, 29]) and turn to the problem of optimizing this trajectory. The nonpositive curvature of \mathcal{S} allows for a time-optimization which does not require explicit construction of \mathcal{S} . We detail an algorithm for transforming any given

edge-path to a time-optimal normal cube path.

From Definition 4.3, an n -dimensional cube C of \mathcal{S} can be represented by a set of n commutative actions $\{\Phi_i\}$ along with an admissible state U . In the following algorithm, we suppress the state for notational convenience and consider C as the set of actions. Addition and subtraction is defined by adding or taking away admissible commutative actions to the list.

ALGORITHM 8.1 (TimeGeodesic). Given: a cube path $\mathcal{C} = \{C_i\}_{i=1..N}$ in \mathcal{S} .

- 1: Let $N := \text{Length}(\mathcal{C})$.
- 2: Call $\text{ShrinkCubePath}(\mathcal{C})$.
- 3: If $\text{Length}(\mathcal{C}) < N$ then 1: else stop.

ALGORITHM 8.2 (ShrinkCubePath). Given: a cube path $\mathcal{C} = \{C_i\}_{i=1..N}$ in \mathcal{S} .

- 1: Let $i = 1$.
- 2: Let $X := \text{Commute}(C_i; C_{i+1})$.
- 3: Update $C_i := C_i + X$;
- 4: Update $C_{i+1} := C_{i+1} - X$.
- 5: Call $\text{ExciseTrivial}(C_{i+1})$.
- 6: Call $\text{CommonEdge}(C_{i-1}, C_i)$.
- 7: Call $\text{ExciseTrivial}(C_{i-1}; C_i)$.
- 8: If $X = \emptyset$ then $i = i + 1$.
- 9: If $C_{i+1} \neq \emptyset$ then 2: else stop.

SUBROUTINE 8.3 (Commute). Given: a pair of cubes $C_j = \{\Phi_{\alpha_i}^j\}_{i=1..m}$ and $C_{j+1} = \{\Phi_{\beta_i}^{j+1}\}_{i=1..n}$,

- 1: Let $S := \bigcup_i \text{SUP}(\Phi_{\alpha_i}^j)$
- 2: Let $T := \bigcup_i \text{TR}(\Phi_{\alpha_i}^j)$
- 3: For $i = 1..n$ return $\Phi_{\beta_i}^{j+1}$ if
 - 3.1: $S \cap \text{TR}(\Phi_{\beta_i}^{j+1}) = \emptyset$; and
 - 3.2: $T \cap \text{SUP}(\Phi_{\beta_i}^{j+1}) = \emptyset$.

Subroutine **ExciseTrivial** checks a cube for an empty list, removes these from the path, and re-indexes the cube path, thus reducing the length. Subroutine **CommonEdge** checks a cube against adjacent cubes in the sequence for common edges and deletes these, returning a genuine cube path (recall Definition 7.2).

Subroutine **Commute** takes as its argument a pair of cubes C_i and C_{i+1} and returns those elements of C_{i+1} which commute with *all* elements of C_i . The following lemma is a manipulation of the definitions.

LEMMA 8.4. *The result of $\text{Commute}(C_i; C_{i+1})$ is precisely the set of edges in $\text{St}(C_i) \cap C_{i+1}$.*

THEOREM 8.5. *Given a cube path \mathcal{C} , Algorithm 8.1 computes a globally time-optimal path in the homotopy class of \mathcal{C} .*

PROOF: Recall that in this setting, the length of a cube path refers to the number of cubes in the path; this equals the time required to execute the reconfiguration.

Repeated calls to Algorithm **ShrinkCubePath** eventually leave a cube path fixed. To

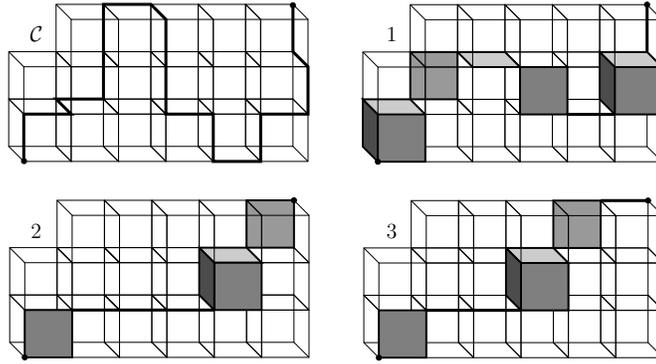


FIG. 8.1. Three rounds of shortening an edge path \mathcal{C} . Cube path length stabilizes after two calls to `ShrinkCubePath`. Three more calls produces a normal cube path.

show this, note that the integer-valued function $f(\mathcal{C}) := \sum_i i \cdot \dim(C_i)$ decreases in each call to `ShrinkCubePath` which changes the path. Lemma 8.4 implies that `ShrinkCubePath` leaves a cube path fixed if and only if it is a normal cube path. From Theorem 7.3, this yields a time-optimal path.

However, Algorithm `TimeGeodesic` calls `ShrinkCubePath` only until the *length* of the cube path is unchanged. Thus it remains to show that once the length of a cube path is unchanged by a call to `ShrinkCubePath`, this is equal to the length of the associated normal cube path. To show this, assume that `ShrinkCubePath` stops without excising any trivial cubes (i.e., shortening the length). We show that further calls merely redistribute actions earlier in the list \mathcal{C} without deleting any permanently. Without a loss of generality, assume that the first call to `ShrinkCubePath` does not change the length of the cube path, but that in the subsequent call, the cube C_{i+1} is eliminated via C_i commuting with a portion of C_{i-1} : see the schematic of Fig. 8.2.

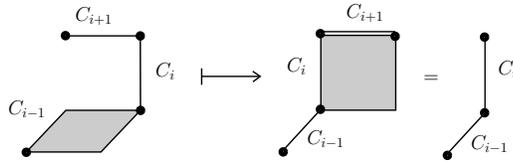


FIG. 8.2. Schematic of exchanging actions in sequential rounds. Each dimension in this cartoon represents a multi-dimensional cube.

Since C_{i+1} is not eliminated in the first call, there must have been actions in C_i which did not commute with those of C_{i-1} but which are subsequently pushed backward in the second call. In addition, all the actions of C_{i+1} are contained in (“parallel to” in the figure) the subset of C_{i-1} which commutes with C_i in the second call. This being the case, the commutativity of C_{i+1} with C_i would have occurred in the first call: contradiction. \diamond

It is not surprising that Algorithm 8.1 converges to a *locally* time-minimal solution. The interesting implication is that the cube-path obtained is the *global* minimum for all paths obtainable from the initial: there is no way to get a quicker reconfiguration path by first

lengthening the path and then shortening. This is the boon of nonpositive curvature.

In Fig. 8.1, we illustrate the effect of Algorithm 8.1 on an initial edge-path in a (very simple Euclidean) state complex. Here, the two types of path-improvement are displayed: (1) commutative moves are executed simultaneously, resulting in higher dimensional cubes being utilized; (2) redundant “backtracking” is eliminated.

EXAMPLE 8.6. Consider a decentralized shape-planning algorithm which takes any connected aggregate of planar hex cells and simplifies it (using the catalogue of Fig. 3.1[right]) to a straight line by moving cells one at a time. This yields a shape-planner by composing a path from the initial shape to the straight line with the inverse of the path from the final shape to the straight line. See Fig. 8.3[top] for a path which translates a small “pile” of three hexes across a long fixed base of N hexes in this manner. This type of reconfiguration path requires precisely $3N + 13$ edges (or time-steps). Using this edge-path as an input to Algorithm 8.1 yields a path of length $N + 2$: as illustrated in Fig. 8.3[bottom], commutative motions are performed simultaneously and backtracking is eliminated. The speed-up in elapsed time from $3N + 13$ to $N + 2$ is, roughly, a factor of three — the “average” dimension of the state complex near the initial edge path. This corresponds to the fact that most of the steps involve simultaneous motion of all three hexes.

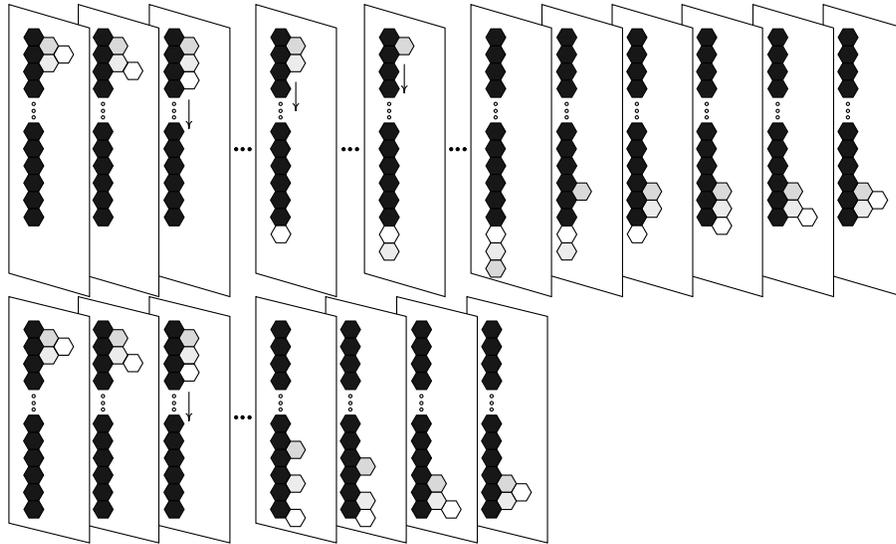


FIG. 8.3. An example of a planar hex system reconfiguration under the action of Algorithm 8.1: [top] an unsimplified edge path; [bottom] the simplified reconfiguration path. The horizontal coordinate denotes time. Dark hexes correspond to fixed cells.

The complexity of Algorithm 8.1 depends on both the catalogue of generators for the system and the length N of the input edge path. However, the dependence on the catalogue arises only in Subroutine `Commute`, which tests various pairs of subsets of the workspace for disjointness; the sizes of these subsets are governed by the generators. Assuming a fixed catalogue, then, there is a constant bound on the time required for a single call to `Commute`, and we may analyze the complexity of Algorithm 8.1 as a function of N alone.

Notice that Algorithm 8.1 calls `ShrinkCubePath` at most N times, since it stops as soon as the length of the cube path fails to decrease. Each time an individual loop of `ShrinkCubePath` is executed, the quantity $\sum \dim(C_j)$ decreases. Since this quantity equals N initially, `ShrinkCubePath` calls the subroutines `Commute`, `ExciseTrivial`, and `CommonEdge` each at most N times. Each of these has a constant running time, with `Commute` being the only one depending on the particulars of the metamorphic system. Therefore the complexity of the entire Algorithm 8.1 is $O(N^2)$, with the constant determined by the running time of `Commute`.

The worst case for Algorithm 8.1 seems to be achieved in a totally flat 2-d state complex by a path consisting of two perpendicular segments of length $N/2$. (This case requires $N^2/2$ runs of `ShrinkCubePath`.) The presence of true negative curvature in a state complex greatly increases the speed of convergence.

9. Discussion. Our principal contributions are:

1. A *mathematical definition* of a metamorphic system which encompass many models currently studied and suggests seemingly unrelated (and often simpler) systems. Given the difficulty of building large metamorphic systems, simpler examples possessing the same formal structure may be valuable.
2. *The state complex*, whose naturality is manifested on the level of its topology (\mathcal{S} can be homotopically simple) and its geometry (non-positive curvature is universal, mathematically helpful, and highly desirable for proving theorems).
3. *An algorithm* for path-optimization which does not require construction of the entire state complex, and thus is compatible with decentralized planners.

There are drawbacks to this approach. Primary among them are the dual dilemmas that shape planning is inherently complex, and that there are many types of reconfiguration possible. A state complex approach is not meant for all systems. Indeed, it is possible to design degenerate metamorphic systems with little to no commutativity. Nevertheless, paying attention to the geometry lurking behind our higher-dimensional versions of transition graphs leads to a nontrivial result on the optimality of path-shortening.

It remains an important computational question to determine whether the homotopy class of the initial path (given, say, from a distributed algorithm) is optimal in the sense that its geodesic is the shortest among all homotopy classes. This problem becomes more difficult in the presence of negative curvature: a breadth-first search in the state complex implicates volumes which grow exponentially with the radius. Interestingly, whereas our algorithm runs better in the presence of more negative curvature, searching for the best homotopy class gets easier as the complex gets flatter.

Finally, in this initial work, we have focused only on the first-order problem of shape planning. We hope that the mathematical framework here suggested finds uses in more sophisticated task-planning problems for metamorphic and reconfigurable robots.

Appendix A. Shape complexes.

State complexes, despite their relative simplicity over the transition graph, nevertheless typically contain a large number of cells. For lattice-based systems, and in some other cases as well, a smaller complex called the *shape complex* can be used to encode all of the local reconfigurations, without losing any essential information from the state complex. The idea is to exploit the symmetries of the domain.

Consider the (lattice-based) pivoting hex system of Example 3.1. The underlying lattice has translational symmetry; for instance, in the idealized case where the workspace is the entire (infinite) lattice with no obstacles, the state complex \mathcal{S} shares these symmetries. If we take a quotient of \mathcal{S} by the actions of these translations, we obtain the shape complex $\mathcal{S}h$. (The precise definition follows shortly.)

Observe that $\mathcal{S}h$ is much smaller than \mathcal{S} ; indeed, even if the workspace is infinite, the complex $\mathcal{S}h$ is compact (provided the system consists of connected sets of, say, N cells). Yet, $\mathcal{S}h$ carries substantially the same information as \mathcal{S} . Whereas \mathcal{S} keeps track of both the shape and the location of the aggregate, the quotient $\mathcal{S}h$ ignores the location.

Note that no information about the system is lost in the process of forming the quotient; \mathcal{S} can be completely reconstructed from $\mathcal{S}h$. Indeed, \mathcal{S} is a *covering space* of $\mathcal{S}h$ (with covering transformation group equal to the group of symmetries of the lattice). In particular $\mathcal{S}h$ and \mathcal{S} share the same universal covering space, so $\mathcal{S}h$ (just like \mathcal{S}) is non-positively curved.

Here is the definition of the shape complex; compare with Definition 4.3.

DEFINITION A.1. The **shape complex** $\mathcal{S}h$ of a (lattice-based) local reconfigurable system is the abstract cube complex whose k -dimensional cubes are equivalence classes $[U; (\Phi_{\alpha_i})_{i=1}^k]$ where

1. $(\Phi_{\alpha_i})_{i=1}^k$ is a k -tuple of commuting actions of generators ϕ_{α_i} ;
2. U is some state for which all the actions $(\Phi_{\alpha_i})_{i=1}^k$ are admissible; and
3. $[U_0; (\Phi_{\alpha_i})_{i=1}^k] = [U_1; (\Phi_{\beta_i})_{i=1}^k]$ if and only if there is a translation $T : \mathcal{L} \rightarrow \mathcal{L}$ such that the list $(\Phi_{\beta_i} \circ T)$ is a permutation of (Φ_{α_i}) , and $U_0 = U_1 \circ T$ on the set $\mathcal{W} - \bigcup_i \Phi_{\alpha_i}(\text{SUP}(\phi_{\alpha_i}))$.

The boundary of a k -cube is obtained just as it is in \mathcal{S} .

EXAMPLE A.2. The shape complex for the system of Example 3.1 (with the generator of Fig. 3.1[left]) having a total of $N = 3$ occupied cells is shown in Fig. A.1. There are nine square cells. Two pairs of edges are identified according to the matching arrows, yielding a Möbius strip; additionally, the three black vertices are identified to one, as are the three white vertices. The shape complex is therefore a Möbius strip with some of the boundary points identified.

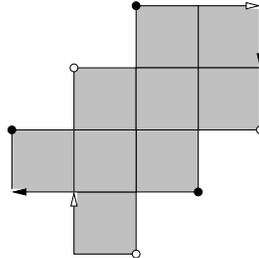


FIG. A.1. The shape complex for a pivoting hex system consisting of three cells.

Macros. The most obvious advantage of the shape complex is that the discovery of a single shortest path in $\mathcal{S}h$ simultaneously solves several state-changing problems in the complex \mathcal{S} . This is particularly helpful in systems which require frequent transportation

within the workspace (perhaps to perform various tasks in different areas). For instance, suppose the workspace includes several narrow corridors through which the robot must travel at various times. At each encounter with such a corridor, it must change to a shape which is skinny enough to navigate the passageway. But this is the same problem at each corridor, even though it occurs in different places in the state complex: the three hexagons in Example A.2 can move from a triangular shape to a linear shape by the same sequence of moves, wherever the problem may arise in \mathcal{L} .

In the shape complex, we view this as a single shortest-path problem, whose solution we may *lift* to various locations in the domain. In this sense a shortest path P in \mathcal{Sh} can be viewed as a macro for solving state-changing problems in \mathcal{S} .

A drawback to the shape complex is that it neglects the shape of the workspace \mathcal{W} . In the presence of obstacles or at the edge of \mathcal{W} , a path P in \mathcal{Sh} may fail to lift to a path in \mathcal{S} . Thus before lifting a path P one must check that the actions are admissible at the appropriate locations in \mathcal{W} .

REFERENCES

- [1] A. Abrams, *Configuration spaces and braid groups of graphs*. Ph.D. thesis, UC Berkeley, 2000.
- [2] A. Abrams and R. Ghrist, *Finding topology in a factory: configuration spaces*, Amer. Math. Monthly (109), 140–150, 2002.
- [3] A. Abrams and R. Ghrist, *Cubical complexes for reconfigurable systems*, in preparation.
- [4] M. Bridson and A. Haefliger, *Metric Spaces of Nonpositive Curvature*, Springer-Verlag, Berlin, 1999.
- [5] Z. Butler, S. Byrnes, and D. Rus, *Distributed motion planning for modular robots with unit-compressible modules*, in Proc. IROS 2001.
- [6] Z. Butler, K. Kotay, D. Rus, and K. Tomita, *Cellular automata for decentralized control of self-reconfigurable robots*, in Proc. IEEE ICRA Workshop on Modular Robots, 2001.
- [7] A. Castano, W.M. Shen, and P. Will, *CONRO: Towards deployable robots with inter-robots metamorphic capabilities*, Autonomous Robots **8**(3): 309-324, 2000.
- [8] G. Chirikjian, *Kinematics of a metamorphic robotic system*, in Proc. IEEE ICRA, 1994.
- [9] G. Chirikjian and A. Pamecha, *Bounds for self-reconfiguration of metamorphic robots*, in Proc. IEEE ICRA, 1996.
- [10] G. Chirikjian, A. Pamecha, and I. Ebert-Uphoff, *Evaluating efficiency of self-reconfiguration in a class of modular robots*, J. Robotics Systems **13**(5): 317-338, 1996.
- [11] D. Epstein et al., *Word Processing in Groups*. Jones & Bartlett Publishers, Boston MA, 1992.
- [12] R. Ghrist, *Configuration spaces and braid groups on graphs in robotics*, AMS/IP Studies in Mathematics volume 24, 29-40, 2001.
- [13] R.W. Ghrist and D.E. Koditschek, *Safe cooperative robot dynamics on graphs*, SIAM J. Cont. & Opt. **40**(5), 1556–1575, 2002.
- [14] M. Gromov, *Hyperbolic groups*, in Essays in Group Theory, MSRI Publ. **8**, Springer-Verlag, 1987.
- [15] K. Kotay and D. Rus, *The self-reconfiguring robotic molecule: design and control algorithms*, in Proc. WAFR, 1998.
- [16] C. McGray and D. Rus, *Self-reconfigurable molecule robots as 3-d metamorphic robots*, in Proc. Intl. Conf. Intelligent Robots & Design, 2000.
- [17] S. Murata, H. Kurokawa, and S. Kokaji, *Self-assembling machine*, in Proc. IEEE ICRA, 1994.
- [18] S. Murata, H. Kurokawa, E. Yoshida, K. Tomita, and S. Kokaji, *A 3-d self-reconfigurable structure*, in Proc. IEEE ICRA, 1998.
- [19] S. Murata, E. Yoshida, A. Kamikura, H. Kurokawa, K. Tomita, and S. Kokaji, *M-TRAN: Self-reconfigurable modular robotic system*, IEEE-ASME Trans. on Mechatronics **7**(4), 431-441, 2002.
- [20] A. Nguyen, L. Guibas, and M. Yim, *Controlled module density helps reconfiguration planning*, in Proc. WAFR, 2000.
- [21] G.A. Niblo and L.D. Reeves, *The geometry of cube complexes and the complexity of their fundamental groups*, Topology, **37**(3), 621-633, 1998.
- [22] A. Pamecha, I. Ebert-Uphoff, and G.S. Chirikjian, *Useful metric for modular robot motion planning*, IEEE Trans. Robotics & Automation, **13**(4), 531-545, 1997.
- [23] D. Rus and M. Vona, *Crystalline robots: Self-reconfiguration with unit-compressible modules*, Au-

- tonomous Robots, **10**(1), 107-124, 2001.
- [24] J. Walter, J. Welch, and N. Amato, *Distributed reconfiguration of metamorphic robot chains*, in Proc. ACM Symp. on Distributed Computing, 2000.
 - [25] J. Walter, E. Tsai, and N. Amato, *Choosing good paths for fast distributed reconfiguration of hexagonal metamorphic robots*, In Proc. IEEE ICRA, 2002.
 - [26] J.E. Walter, J.L. Welch, and N.M. Amato, *Concurrent metamorphosis of hexagonal robot chains into simple connected configurations*, IEEE Trans. Robotics & Automation **15**(6), 1035-1045, 1999.
 - [27] M. Yim, *A reconfigurable robot with many modes of locomotion*, in Proc. Intl. Conf. Adv. Mechatronics, 1993.
 - [28] M. Yim, J. Lamping, E. Mao, and J. Chase, *Rhombic dodecahedron shape for self-assembling robots*, Xerox PARC Tech. Rept. P9710777, 1997.
 - [29] M. Yim, Y. Zhang, J. Lamping, and E. Mao, *Distributed control for 3-d metamorphosis*, Autonomous Robots J., 41-56, 2001.
 - [30] E. Yoshida, S. Murata, K. Tomita, H. Kurokawa, and S. Kokaji, *Distributed formation control of a modular mechanical system*, in Proc. Intl. Conf. Intelligent Robots & Sys., 1997.