# Shape Complexes for Metamorhpic Robots*

Robert Ghrist

University of Illinois, Urbana, IL 61801, USA

**Abstract.** A **metamorphic** robotic system is an aggregate of identical robot units which can individually detach and reattach in such a way as to change the global shape of the system. We introduce a mathematical framework for defining and analyzing general metamorphic systems. This formal structure combined with ideas from geometric group theory leads to a natural extension of a configuration space for metamorphic systems — the **shape complex** — which is especially adapted to parallelization. We present an algorithm for optimizing reconfiguration sequences with respect to elapsed time. A universal geometric property of shape complexes — **non-positive curvature** — is the key to proving convergence to the globally time-optimal solution.

## 1 Introduction

In recent years, several groups in the robotics community have been modeling and building **reconfigurable** or, more specifically, **metamorphic systems** (*e.g.*, [6,11,13,19,20]). 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. There are as many models for such systems as there are researchers in the sub-field: 2-d and 3-d lattices; hexagonal, square, and dodecahedral cells; pivoting or sliding motion [6,11–15,19,20]. The common features of these systems are an aggregate of lattice-based cells having prescribed transitions from one shape to another.

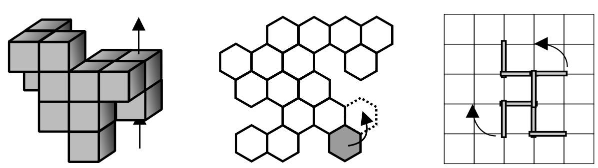

**Fig. 1.** Metamorphic systems may be built on a variety of lattice structures with sliding or pivoting motion.

The primary problem to be solved in such systems is the shape-planning problem: how to move from one shape to another via legal moves. One centralized approach [7,17] is to build a **transition graph** whose vertices are the

* Supported in part by National Science Foundation Grant DMS-0134408.

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.

We propose to extend the notion of a configuration space to metamorphic systems. The idea: consider the transition graph described above as the 1-skeleton of a **shape complex**. Assume that from a given shape there are two possible moves which are physically independent (or, more suggestively, "commutative"): *i.e.*, these moves can be executed simultaneously. In 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 $k$-cubes corresponding to $k$ commutative motions. The result is a cubical complex which has several advantages over the transition graph:

1. **Size.** The shape complex is often simpler than the transition graph: *e.g.*, the 1-d graph of an $n$-dimensional cube has $n2^{n-1}$ edges. This figure belies the simplicity of the single cube.
2. **Time.** Geodesics on this complex cut across the diagonals of cubes whenever possible: one performs all possible commutative motions simultaneously, yielding a speed-up by a factor of the dimension of the complex.
3. **Shape.** The global geometry/topology of the shape complex carries information about the metamorphic system. For certain examples, the topology of the shape complex "converges" upon refining the lattice. In addition, only special geometries can be realized as the shape complex of a local metamorphic system: commutativity in reconfiguration leads to an abhorrence of positive curvature in the shape complex.

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

For large systems, the problem of computing the shape complex and designing geodesics in order to perform shape planning is computationally infeasible for two reasons: (1) the size of the complex is often exponential in the number of robot cells; and (2) any control scheme induced by geodesic construction is necessarily centralized. Several researchers have begun building decentralized control algorithms for shape planning [4,5,18,21,22]. Such algorithms have the advantage of speed; however, the reconfiguration paths are rarely if ever optimal.

As an application of our techniques, we present in Sections 7-8 an algorithm for trajectory optimization which takes as its argument an arbitrary edge path in the transition graph. Algorithm 5 then performs a type of **curve shortening** within the shape complex. A deep theorem about the curvature of all shape complexes (Theorem 1) 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 systems which employ continuous reconfiguration for locomotive gaits are not covered by our definitions. We note, however, that certain locomotive reconfigurable systems can be thought of as lattice-based tiles by amalgamating subsystems [12]. In addition, our definitions can easily be extended to general non-lattice reconfigurable systems [2].

## 2 A mathematical definition

While it is easy to generate examples of what is meant by a metamorphic system, 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.

A local metamorphic system is a collection of states on a lattice, where each state is thought of as an indicator function for the aggregate. 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 domain. The adjective "local" refers to legality criteria: anywhere in the domain 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 1.** Let $\mathcal{L}$ denote a lattice in $\mathbb{R}^k$ and $\mathcal{D} \subset \mathcal{L}$ be some domain. The **catalogue** $\mathsf{C}$ for a local metamorphic system on $\mathcal{D}$ is a collection of **generators**. Each generator $\phi \in \mathsf{C}$ consists of (1) the **support**, $\sup(\phi) \subset \mathcal{L}$; (2) the **trace** of the move, $\mathrm{tr}(\phi) \subset \sup(\phi)$; and (3) an unordered pair of local states $\hat{U}_{0,1} : \sup(\phi) \to \{0,1\}$ satisfying[1]

$$\hat{U}_0\Big|_{\sup(\phi)-\mathrm{tr}(\phi)} = \hat{U}_1\Big|_{\sup(\phi)-\mathrm{tr}(\phi)}. \tag{1}$$

Otherwise said, the local states are equal on $\sup(\phi) - \mathrm{tr}(\phi)$.

**Definition 2.** An **action** of a generator $\phi \in \mathsf{C}$ is a rigid translation $\Phi : \sup(\phi) \hookrightarrow \mathcal{D}$. Given a state $U : \mathcal{D} \to \{0,1\}$, an action $\Phi$ is said to be **admissible** at $U$ if $\hat{U}_0 = U \circ \Phi$. In this case, we write

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

**Definition 3.** A **local metamorphic system** on $\mathcal{D}$ is a collection of states $\{U_\alpha : \mathcal{D} \to \{0,1\}\}$ closed under all possible admissible actions of generators in the catalogue $\mathsf{C}$.

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

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

A metamorphic system with **obstacles** $\mathcal{O} \subset \mathcal{D}$ satisfies in addition the property that $U|_{\mathcal{O}} \equiv 1$ for all states $U$, and that for each admissible action of each $\phi \in \mathsf{C}$,

$$\Phi(\mathrm{tr}(\phi)) \cap \mathcal{O} = \emptyset. \tag{2}$$

Of course, the set $\mathcal{O}$ represents those portions of the domain $\mathcal{D}$ on which agents are forbidden to occupy or to pass through. They do count as "occupied" 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 $\mathcal{O}$).

## 3 Examples

*Example 1.* [2-d hex with pivots] This reconfigurable system has a simple catalogue of six generators (or one, up to rigid rotations), represented in Fig. 2. The catalogue is chosen so that the aggregate does not change its topology but only its shape (a slightly smaller support allows for topology changes). To model a fixed "base" cell (which is, say, affixed to a power source as in [15]), one establishes this cell as an obstacle $\mathcal{O}$.

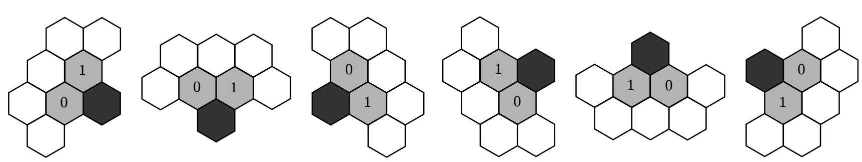


**Fig. 2.** The catalogue for a 2-d hexagonal lattice system with pivots. Black cells are occupied, white are unoccupied, and grey cells denote the trace. Cells turned on in local state $\hat{U}_0$ and $\hat{U}_1$ are indicated numerically.

*Example 2.* [2-d square lattice] In Fig. 3, we display a generator for a planar system in which rows and columns [not pictured] of an aggregate of square cells can slide. There are in fact several generators represented in "shorthand." The four white squares marked with dots can be turned on or off — but if on, their immediate neighbor (indicated by an arrow) must also be turned on. 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 minus the two endpoints.

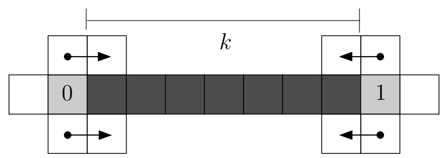


**Fig. 3.** One set of generators for a sliding-squares system.

*Example 3.* [2-d articulated planar arm] Consider as a domain $\mathcal{D}$ the set of edges in the planar integer lattice. The catalogue consists of two generators, pictured in Fig. 4. Beginning with a state having $N$ vertical edges end-to-end (with the bottom one fixed as an obstacle), the metamorphic system thus generated models the position of an articulated robotic arm which can (1) rotate at the end and (2) flip corners as per the diagram.

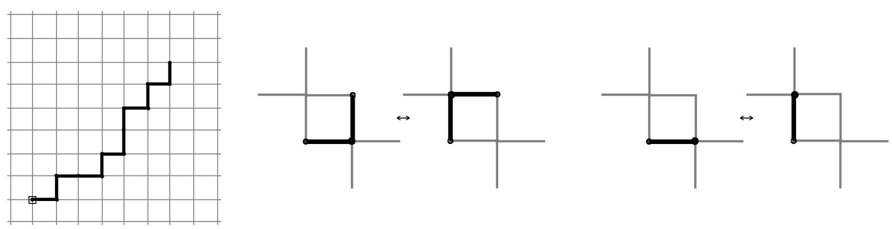

**Fig. 4.** A positive articulated robot arm example [left] with fixed endpoint. One generator [center] flips corners and has as its trace the central four edges. The other generator [right] rotates the end of the arm, and has trace equal to the two activated edges.

If rotations of these generator are not included, the robot arm is **positive** in the sense that the arm may extend up and to the right only. More general versions with multiple interacting arms can be achieved by adding an “attach-detach” generator to the catalogue.

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: the following example is particularly interesting.

*Example 4.* Consider a graph $\Gamma$ in the Euclidean lattice (all edges are edges in the Euclidean lattice in $\mathbb{R}^k$). The catalogue consists of a single generator whose support and trace is precisely the closure of an edge in this lattice. 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 metamorphic system generated from a state $U$ on $\Gamma$ with $N$ vertices evaluating to 1 mimics an ensemble of $N$ non-colliding points on $\Gamma$.

More abstract examples of metamorphic examples include the space of triangulations of convex polygons with edge-flipping as the generator, examples arising from word representations in group theory, and certain multi-step assembly processes [2].

## 4 The shape complex

In the robotics literature, one often models a configuration space for a metamorphic system with a transition graph which coordinates 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: [7,17]). Our departure is to make the transition graph the 1-skeleton of a cubical complex (the analogue of a simplicial complex made out of abstract cubes) which coordinates "commutative" motions.

**Definition 4.** 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

$$\Phi_{\alpha_i}(\mathrm{tr}(\phi_{\alpha_i})) \cap \Phi_{\alpha_j}(\sup(\phi_{\alpha_j})) = \emptyset \quad \forall i \neq j. \tag{3}$$

*Example 5.* 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. 5 [left]. Compare this with a planar sliding block example as illustrated in Fig. 5

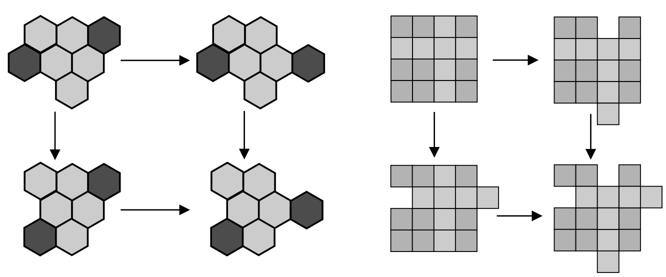

**Fig. 5.** Examples of commuting [left] and noncommuting [right] actions in planar systems.

[right]. Although the pair of moves illustrated forms a square[2] in the transition graph, this particular pair of actions does not commute. Physically, it is

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

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 shape complex has an abstract $k$-cube for each collection of $k$ admissible commuting actions:

**Definition 5.** The **shape 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 $\phi_{\alpha_i}$;
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 $(\beta_i)$ is a permutation of $(\alpha_i)$ and $U_0 = U_1$ on the set $\mathcal{D} - \bigcup_i \Phi_{\alpha_i}(\sup(\phi_{\alpha_i}))$ .

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,

$$\partial[U; (\Phi_{\alpha_i})_{i=1}^k] = \bigcup_{j=1}^{k} \left([U; (\Phi_{\alpha_j})_{j\neq i}] \cup [\phi_{\alpha_i}[U]; (\Phi_{\alpha_j})_{j\neq i}]\right) \tag{4}$$

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

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

(b) A 1-cell of $\mathcal{S}$ is an equivalence class of the form $[U; (\Phi)]$. 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 $[\phi[U]; (\cdot)]$. □

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

*Example 6.* Consider the 2-d square lattice row/column sliding system whose catalogue is illustrated in Fig. 3. If we consider a system with obstacles in the form of a $p$-by-$q$ rectangle generated from the state of Fig. 6 [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 shape complex is that of Fig. 6 [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.

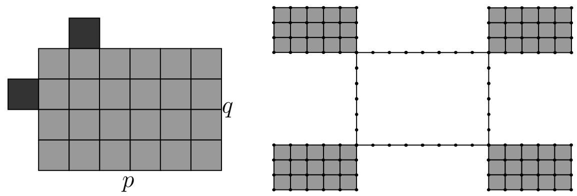


**Fig. 6.** The planar sliding square example [left] with two movable blocks [in black] and a $p$-by-$q$ obstacle set [in grey] yields a shape complex $\mathcal{S}$ that is topologically a circle [right].

*Example 7.* The shape complex associated to the positive articulated robot arm of Example 3 in the case $N = 5$ is given in Fig. 7. Note that there can be at most three independent motions (when the arm is in a "staircase" configuration); hence the shape complex has top dimension three. Notice also that although the transition graph for this system is complicated, the shape complex itself is topologically trivial (contractible).

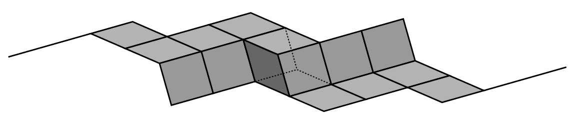

**Fig. 7.** The shape complex of a 5-link positive arm has one cell of dimension three, along with several cells of lower dimension.

*Example 8.* In the system of Example 4 with the graph being a $K_5$ (the complete graph on five vertices) and $N = 2$, the shape complex is a two-dimensional closed orientable surface. A simple combinatorial argument reveals that the Euler characteristic is $-10$, implying that the shape complex has genus six.

## 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 section argues that completing the transition graph to the shape complex is "natural" — the shape complex simplifies the transition graph and endows it with 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 cycles) blows up as well. This is not always so with the shape complex: in several key examples, the topology of $\mathcal{S}$ converges for large systems.

*Example 9.* Recall the shape complex associated to the metamorphic system of $N$ points on a graph $\Gamma$, Example 4. Consider a refinement of the underlying lattice of $\Gamma$ which inserts additional vertices along edges. It follows from the techniques of [1] that the shape complex of this refined system has the same topological type (up to homotopy equivalence) after a fixed bound on the refinement ($N-2$ additional vertices per edge). Furthermore, this "stabilized" shape complex is in fact homotopic to the topological configuration space of $N$ points of $\Gamma$ (what one expects as the number of refinements goes to infinity).

*Example 10.* Recall the articulated robot arm of Example 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). It can be shown [2] that the shape complex of the refined system does not change homotopy type, even though the dimension of the complex goes up by a factor of two. In the limit as the number of refinements goes to infinity, the shape complex approximates the configuration space of a smooth curve of fixed length.

Not all systems obey this stabilization principle. For example, in the hex-lattice systems of Example 1, the shape complex possesses loops represented by "braiding" of pairs of hex elements about the boundary of the aggregate. Since a lattice rescaling increases the number of cells along the boundary, such loops grow in number. However, if we scale the geometry of the shape complex in accordance with the lattice scaling, these loops shrink to length zero in the continuum limit.

## 6 The geometry of $\mathcal{S}$

Shape complexes are built from Euclidean (i.e., flat) cells of unit length. However, this does not imply that the complex is, as a whole, flat. Indeed, non-zero curvature can be concentrated at places where several cells meet. A simple example appears in Fig. 8: 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. Such an extension of curvature to general metric spaces is made precise in Gromov's work on curved metric spaces [10] (extending the classical work of Andronov, 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

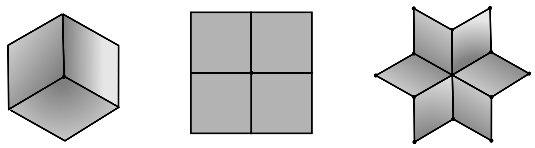

**Fig. 8.** Curvature about a vertex: $+, 0, -$.

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 associated chord.

**Definition 6.** 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 $d \leq d'$.

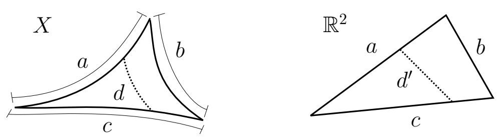


**Fig. 9.** Comparison triangles measure curvature bounds.

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 spaces: for example, the universal cover is contractible.

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

**Theorem 1 ([2]).** *The shape complex $\mathcal{S}$ of any local metamorphic system is nonpositively curved.*

The proof requires some machinery which we do not include here. The key observations required are that (1) commutativity of a set of actions is independent of the states implicated; and (2) any collection of pairwise commutative actions is totally commutative. These conditions, properly translated into the language of Gromov's theory, imply that the complex never has positive curvature.

*Example 11.* To see where the NPC property can fail, consider the following non-local metamorphic system. Fig. 10[left] gives a generator for a planar hexagonal pivoting system which (along with its rotations) allows for local disconnection of the aggregate. If we impose a global rule that requires the aggregate to be connected, for, say, considerations of power transmission, then we may create positive curvature. Fig. 10 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.[3]

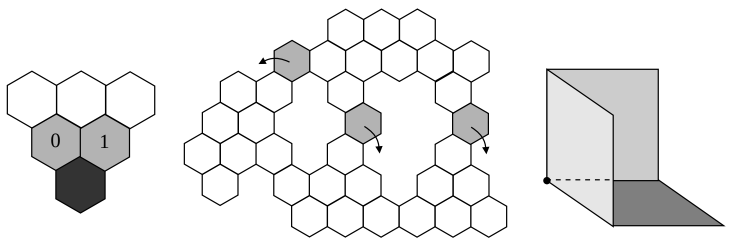


**Fig. 10.** A generator which allows for local disconnection [left] admits configurations [center] for which pairs of actions leave the aggregate connected, but triples do not; this non-local system has positive curvature in the shape complex at this state [right].

## 7 Geodesics and time on $\mathcal{S}$

Besides displaying a unifying geometric feature, the nonpositive curvature has implications in the shape planning problem.

**Corollary 2** *Geodesics on any shape complex are unique in their homotopy class.*

PROOF: Every NPC space has contractible universal cover [3]. □

Spheres, for example, do not share this property: there are many geodesics between, say, antipodal points which are homotopic. This corollary is the key ingredient in the applications of NPC geometry to path-planning on a configuration space.

One expects geodesics on $\mathcal{S}$ to coincide with optimal solutions to the shape planning problem. Corollary 2 implies that any length-decreasing algorithm

[3] The shape complex will not be two-dimensional here, but will be locally a product of this with another cubical complex. The positive curvature persists.

*must* converge to the minimal-length curve possible: there are no spurious local minima (or even saddles) to prevent convergence. This good news is tempered with ill: it is not necessarily easy to find the true geodesic quickly.

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 shape complex. Fig. 11[left] illustrates the matter concisely. Consider a portion of a shape 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.

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. 11[left] "slows down" some of the moves unnecessarily in order to maintain the constant slope.

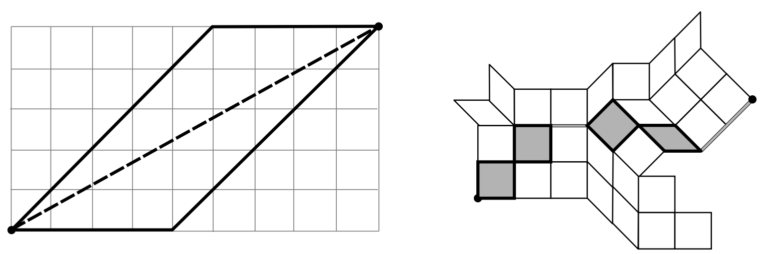

**Fig. 11.** [left] The true geodesic lies within an envelope of time-minimizing paths. No minimal edge-paths are time-minimizing; [right] a normal cube-path in an NPC complex (after [16]).

If, instead of the Euclidean metric on the cells of $\mathcal{S}$, one measures the $\ell^\infty$ norm of tangent vectors (the sup of the components), then, geodesics in this category are paths which are *time minimizing*. By using the techniques of [16], one can prove that **cube paths** are the appropriate class of generalized paths.

**Definition 7.** A **cube path** from vertices $v_0$ to $v_N$ is an ordered sequence of closed cubes $\{C_i\}_1^N$ in $\mathcal{S}$ which satisfy (1) $C_i \cap C_{i+1} = v_i$; 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 $\forall i$ (3) $C_{i+1} \cap \mathsf{St}(C_i) = v_i$, where $\mathsf{St}(C_i)$ is the **star** of $C_i$ (all cells, including $C_i$, which have $C_i$ as a face).

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

**Theorem 3.** *In any shape 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 is a globally time-minimal path over all paths in its homotopy class.*

Existence and uniqueness follows from Theorem 1 and a result of [16] (or, see the constructive algorithm of Section 8 below). Time-minimization follows from geodesic-like properties of normal cube paths: see [2] for proof details.

## 8 Optimizing paths

In cases where the shape complex is sparse — local principal cells being of high dimension with few neighbors — it is possible to compress the size of the transition graph significantly. However, shape-planning via constructing all of $\mathcal{S}$ and determining geodesics is, in general, computationally infeasible: the total size of the shape complex is often exponential in the number of cells in the aggregate. We therefore assume that some shape trajectory has been determined (by a perhaps ad hoc method) and turn to the problem of optimizing this trajectory. The nonpositive curvature of $\mathcal{S}$ allows for a time-optimization which requires exploration of small regions within $\mathcal{S}$. We detail an algorithm for transforming any given edge-path to a time-optimal normal cube path.

From Definition 5, 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 4** (`TimeGeodesic`) Given: a cube path $\mathcal{C} = \{C_i\}_{i=1..N}$ in $\mathcal{S}$.
1: Let $N := \texttt{Length}(\mathcal{C})$.
2: Call $\texttt{ShrinkCubePath}(\mathcal{C})$.
3: If $\texttt{Length}(\mathcal{C}) < N$ then 1: else `stop`.

**Algorithm 5** (`ShrinkCubePath`) Given: a cube path $\mathcal{C} = \{C_i\}_{i=1..N}$ in $\mathcal{S}$.
1: Let $i = 1$.
2: Let $X := \texttt{Commute}(C_i; C_{i+1})$.
3: Update $C_i := C_i + X$;
4: Update $C_{i+1} := C_{i+1} - X$.
5: Call $\texttt{ExciseTrivial}(C_{i+1})$.
6: Call $\texttt{CommonEdge}(\mathcal{C})$.
7: If $X = \emptyset$ then $i = i + 1$.
8: If $C_{i+1} \neq \emptyset$ then 2: else `stop`.

**Subroutine 6** (`Commute`) Given: a pair of cubes $C_j = \{\Phi^j_\alpha\}_{1..m}$ and $C_{j+1} = \{\Phi^{j+1}_\alpha\}_{1..n}$,

1: Let $S := \bigcup_\alpha \mathsf{sup}(\Phi^j_\alpha)$
2: Let $T := \bigcup_\alpha \mathsf{tr}(\Phi^j_\alpha)$
3: For $i = 1..n$ return $\Phi^{j+1}_i$ if
 3.1: $S \cap \mathsf{tr}(\Phi^{j+1}_i) = \emptyset$ ; and
 3.2: $T \cap \mathsf{sup}(\Phi^{j+1}_i) = \emptyset$.

Subroutine `ExciseTrivial` checks the cubes for empty lists, removes these from the path, and reindexes the cube path, thus reducing the length. Subroutine `CommonEdge` checks adjacent cubes in the sequence for common edges and deletes these as necessary, returning a genuine cube path (recall Definition 7).

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

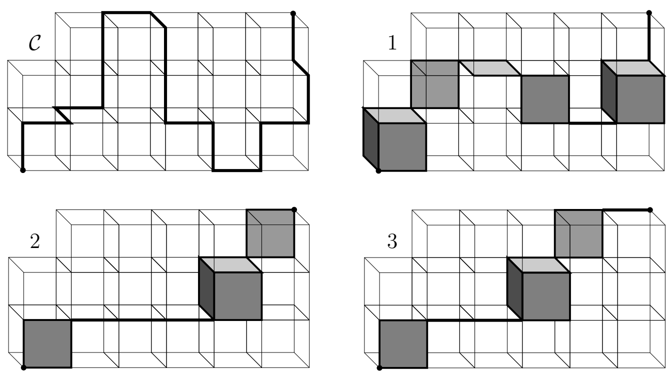


**Fig. 12.** Three rounds of shortening an edge path $\mathcal{C}$. Cube path length stabilizes after two calls to `ShrinkCubePath`. Three more calls produces a normal cube path.

**Theorem 7.** *Algorithm 4 computes a globally time-optimal path in the homotopy class of* $\mathcal{C}$.

PROOF: Repeated calls to Algorithm `ShrinkCubePath` eventually leaves a cube path fixed. To show this, note that the function $f(\mathcal{C}) := \sum_i i \cdot \dim(C_i)$ decreases in each call to `ShrinkCubePath` which changes the path. Lemma 2 implies that Algorithm 5 leaves a cube path fixed if and only if it is a normal cube path. From Theorem 3, this yields a time-optimal path.

The final step is to show that once the length of a cube path in unchanged by a call to `ShrinkCubePath`, this is equal to the length of the associated normal cube path. To show this, assume that `ShrinkCubePath` completes a round without excising any trivial cubes (i.e., shortening the length). We show that further rounds 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. 13.

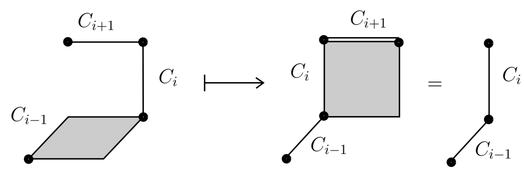


**Fig. 13.** Schematic of exchanging actions in sequential rounds. Each dimension in this cartoon represents a multi-dimensional cube.

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 forward 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. □

It is not surprising that Algorithm 4 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. Fig. 12 gives an example of paths generated by Algorithm 4.

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.

The complexity of Algorithm 4 is best measured as a function of $N$, the number of edges in the initial cube path. In the worst case, Algorithm 4 calls `ShrinkCubePath` $\lceil N/2 \rceil$ times. This worst case is achieved by an initial edge path in a planar shape complex with one corner at the end of edge number $\lfloor N/2 \rfloor$. One can use a 2-stage induction on $N$ to show this (we do not include the full argument for reasons of space). Assume that $N$ is even and that the statement holds for paths of length up to $N$. For a path of length $N+1$, each call to `ShrinkCubePath` does not register the last edge until the first $N$ edges are processed. To keep this path a worst case, the last edge should not commute with any other cell faces of $\mathcal{C}$ until the very last call to `ShrinkCubePath`. Hence, the penultimate vertex remains fixed, and the worst case proceeds from induction: the shape complex is planar and the path has one corner in the middle. The final edge must lie in this plane. The case of $N$ odd follows a slightly different argument: the extra edge is added at the *beginning* of the sequence and commutes at the very last call.

In this worst case example there are $\frac{1}{2}(N^2 - N)$ calls to `Commute`. This can almost certainly be improved with a more sophisticated algorithm. The subroutine `Commute` is an improvement over the more obvious method of computing $\mathtt{St}(C_i)$ at each step — the algorithm merely requires checking if each action of $C_{i+1}$ commutes with every action of $C_i$.

The worst case is a totally flat shape complex (which is desirable in other respects). The presence of true negative curvature in a shape complex greatly increases the speed of convergence. If, as in Example 8, the shape complex has negative curvature at every vertex, then Algorithm 4 converges to the normal cube path with a *linear* number of calls to `Commute`. This transition from quadratic to linear as one passes from zero to negative curvature is a persistent phenomenon [8].

## 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 shape 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, rare, and highly desirable for proving theorems).

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 shape 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.

As an interesting side-note, we observe that in the present context, the optimal reconfiguration path is not always coincident with a geodesic on the configuration space. This principal holds in more generality.

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 a use in more sophisticated task-planning problems for metamorphic and reconfigurable systems.

## References


1. A. Abrams, *Configuration spaces and braid groups of graphs.* Ph.D. thesis, UC Berekeley, 2000.
2. A. Abrams and R. Ghrist, *State complexes for reconfigurable systems*, in preparation.
3. M. Bridson and A. Haefliger, Metric Spaces of Nonpositive Curvature, Springer-Verlag, Berlin, 1999.
4. Z. Butler, S. Byrnes, and D. Rus, *Distributed motion planning for modular robots with unit-compressible modules*, in Proc. IROS 2001.
5. 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.
6. G. Chirikjian, *Kinematics of a metamorphic robotic system*, in Proc. IEEE ICRA, 1994.
7. G. Chirikjian and A. Pamecha, *Bounds for self-reconfiguration of metamorphic robots*, in Proc. IEEE ICRA, 1996.
8. D. Epstein et al., Word Processing in Groups. Jones & Bartlett Publishers, Boston MA, 1992.
9. R. Ghrist, *Configuration spaces and braid groups on graphs in robotics*, AMS/IP Studies in Mathematics volume 24, 29-40, 2001. ArXiv preprint math.GT/9905023.
10. M. Gromov, *Hyperbolic groups*, in Essays in Group Theory, MSRI Publ. **8**, Springer-Verlag, 1987.
11. K. Kotay and D. Rus, *The self-reconfiguring robotic molecule: design and control algorithms*, in Proc. WAFR, 1998.
12. C. McGray and D. Rus, *Self-reconfigurable molecule robots as 3-d metamorphic robots*, in Proc. Intl. Conf. Intelligent Robots & Design, 2000.
13. S. Murata, H. Kurokawa, and S. Kokaji, *Self-assembling machine*, in Proc. IEEE ICRA, 1994.
14. S. Murata, H. Kurokawa, E. Yoshida, K. Tomita, and S. Kokaji, *A 3-d self-reconfigurable structure*, in Proc. IEEE ICRA, 1998.
15. A. Nguyen, L. Guibas, and M. Yim, *Controlled module density helps reconfiguration planning*, in Proc. WAFR, 2000.
16. G. Niblo and L. Reeves, *The geometry of cube complexes and the complexity of their fundamental groups*, Topology, **37**(3), 621-633, 1998.
17. A. Pamecha, I. Ebert-Uphoff, and G. Chirikjian, *Useful metric for modular robot motion planning*, in IEEE Trans. Robotics & Automation, **13**(4), 531-545, 1997.
18. J. Walter, J. Welch, and N. Amato, *Distributed reconfiguration of metamorphic robot chains*, in Proc. ACM Symp. on Distributed Computing, 2000.
19. M. Yim, *A reconfigurable robot with many modes of locomotion*, in Proc. Intl. Conf. Adv. Mechatronics, 1993.
20. M. Yim, J. Lamping, E. Mao, and J. Chase, *Rhombic dodecahedron shape for self-assembling robots*, Xerox PARC Tech. Rept. P9710777, 1997.
21. Y. Zhang, M. Yim, J. Lamping, and E. Mao, *Distributed control for 3-d shape metamorphosis*, Aut. Robots. J., 41-56, 2001.
22. 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.