Pushing Squares Around

Adrian Dumitrescu
Computer Science
University of Wisconsin–Milwaukee
Milwaukee, WI 53201, USA
ad@cs.uwm.edu

János Pach
Courant Institute of Mathematical Sciences
251 Mercer Street
New York, NY 10012-1185, USA
pach@cims.nyu.edu

ABSTRACT

We study dynamic self-reconfiguration of modular metamorphic systems. We guarantee the feasibility of motion planning in a rectangular model consisting of square modules that are allowed to slide along or rotate about one another. That is, we show that any two connected configurations of the same number of modules can be transformed into each other by a sequence of moves so that all intermediate configurations are connected. This settles a conjecture formulated in [6].

Categories and Subject Descriptors

F.2.2 [**Nonnumerical Algorithms**]: Geometric problems and computations—*computations on discrete structures*

General Terms

Algorithms, Theory

Keywords

Motion planning, rectangular metamorphic systems, reconfiguration

1. INTRODUCTION

A modular *metamorphic* system consists of a number of identical modules that can connect to, disconnect from, and relocate relative to adjacent modules (see, for example, [2, 7, 9, 10, 12]). While individual modules are not capable of moving by themselves, the entire system may be able to reconfigure or move to a new position, when its members repeatedly change their positions relative to their neighbors, by rotating or sliding around other modules [1, 7, 11], or by expansion and contraction [10]. It is usually assumed that the entire system must remain connected during reconfiguration.

The *motion planning* problem for such a system is that of computing a sequence of module motions that brings the system from a given initial configuration I into a desired goal configuration F. Depending on the existence of such a sequence of motions, we say that the problem is *feasible* or respectively, *infeasible*.

For instance, in [2, 9] upper and lower bounds on the number of moves needed to change I to F are discussed for a system of hexagonal modules. In [8] it is shown that in such a system any two connected configurations are mutually reachable as long as they do not contain a certain prohibited pattern. Demaine et. al. [3] have considered a family of one-player games, involving the movement of coins from one configuration to another. Moves are restricted so that a coin can only be placed in a free position adjacent to at least two other coins, in contrast to our motion rules that are required to maintain overall connectedness throughout the reconfiguration process.

In Section 2, we present the rectangular model, and in Section 3 we prove that its two motion rules (sliding and rotation) guarantee the feasibility of motion planning for any pair of connected configurations having the same number of modules. This settles a conjecture formulated in [6]. Section 4 presents another rectangular model for which the same property holds. We conclude (Section 5) by discussing several related reconfiguration problems.

2. RECTANGULAR METAMORPHIC SYSTEMS

Consider a plane that is partitioned into a rectangular integer grid of square cells indexed by their center coordinates in the underlying x-y coordinate system. Of the eight *adjacent* cells of cell $c = c_{x,y}$ in the E $(+x)$, W $(-x)$, N $(+y)$, S $(-y)$, NE, SE, NW and SW directions, the four in the E, W, N and S directions are said to be *side-adjacent* to c, while the other four in the NE, SE, NW and SW directions are said to be *corner-adjacent* to c. We denote by $N(c)$ (resp. $NE(c)$) the cell side-adjacent to c in the N direction (resp. the cell corner-adjacent to c in the NE direction). Similar notation is used to denote the cells side-adjacent or corner-adjacent to c in the other axis and diagonal directions.

At any time each cell may be empty or occupied by a module. The *reconfiguration* of a metamorphic system consisting of n modules is a sequence of configurations (distributions) of the modules in the grid at discrete time steps $t = 0, 1, 2, \ldots$, see below. Let V_t be the configuration of the modules at time t, where we often identify V_t with the set of cells occupied by the modules or with the set of their centers. We are only interested in configurations that are

SCG'04, June 8–11, 2004, Brooklyn, New York, USA.

connected, i.e., for each t, the graph $G_t = (V_t, E_t)$ must be connected, where for any t, E_t is the set of edges connecting pairs of cells in V_t that are side-adjacent. V_t *yields* V_{t+1} when one module m moves from its current location to new location in step t. In this paper we restrict ourselves to sequential reconfiguration, in which only one module moves at each discrete time step, as explained above. Note that, according to the above definition, the pattern (set of cells or set of integer points) V_t uniquely determines the edge set E_t so that the graph G_t can be characterized by its vertex set V_t. The union of all closed squares belonging to V_t is a connected point set of area $|V_t|$, which will be denoted by $S(V_t)$.

The following two generic motion rules (Figure 1) define the *rectangular model*. These are to be understood as possible in all axis parallel orientations, in fact generating 16 rules, eight for rotation and eight for sliding.

- *Rotation*: A module m side-adjacent to a stationary module f rotates through an angle of 90° around f either clockwise or counterclockwise. Figure 1(a) shows a clockwise NE rotation. For rotation to take place, both the target cell and the cell at the corresponding corner of f that m passes through (NW in the given example) have to be empty.

- *Sliding*: Let f_1 and f_2 be stationary cells that are side-adjacent. A module m that is side-adjacent to f_1 and adjacent to f_2 slides along the sides of f_1 and f_2 into the cell that is adjacent to f_1 and side-adjacent to f_2. Figure 1(b) shows a sliding move in the E direction. For sliding to take place, the target cell has to be empty.

In order to ensure motion precision, each move is guided by one or two modules that are stationary during the same step. The two motion rules of this model also appear in [5, 6]. A somewhat similar model is presented in [1].

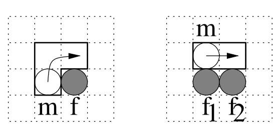

(a) (b)

Figure 1: (a) Clockwise NE rotation and (b) sliding in the E direction. Fixed modules are shaded. The cells in which the moves take place are outlined in the figure.

THEOREM 1. *The set of motion rules of the rectangular model guarantees the feasibility of motion planning for any pair of connected configurations V and V' having the same number of modules. That is, following the above rules, V and V' can always be transformed into each other so that all intermediate configurations are connected.*

We refer to a set of modules that form a straight line chain in the grid, as a *straight chain*. It is easy to construct examples so that neither sliding nor rotation alone can reconfigure them to straight chains. However, in Section 3 we prove that the motion rules of the rectangular model (rotation and sliding, Figure 1) are sufficient to guarantee reachability, while maintaining the system connected at each discrete time step. In [6], it was proved that this is true for a special class of systems, called horizontally convex, where also a distributed algorithm was given for this task in a setting where concurrent moves are allowed.

3. PROOF OF THEOREM 1

3.1 Algorithm outline

Clearly, it is enough to prove that any configuration can be transformed into a straight chain. Then one can append to the sequence of moves that transforms the start configuration into a straight chain, the reversed sequence of moves that transforms the goal configuration into a straight chain to obtain the desired effect. In addition, we make use of the easy fact that the two motion rules permit the relocation of a straight chain from a given initial location to any target location.

Assume without loss of generality that the maximum x-coordinate of a cell in the configuration is 0. Arbitrarily select one of the occupied cells having $x = 0$ as the *base*, say s. The algorithm relocates all the modules except s, one by one, to extend a horizontal chain Z, whose leftmost cell is s. Denote by z the rightmost cell of Z (initially, $s = z$). During the execution of the algorithm, all modules except those in $Z \setminus \{s\}$ have non-positive x coordinates. In each iteration of the algorithm, one module extends the chain in the E direction by one cell. If desired, the algorithm can be easily modified so that reconfiguration is done in a compact space, instead of a straight chain.

3.2 Preliminaries

Sometimes we refer to the cell occupied by a module m by *cell*(m). A pair of modules u, v is said to form a *critical pair* if u and v are corner-adjacent, say at point p, and the other two cells corner-adjacent at point p are empty. See Figure 2. Note that u and v are connected by a path (not shown in the figure). For example, c, d in Figure 4 form a critical pair.

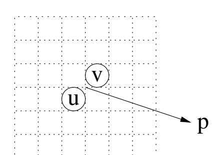

Figure 2: A critical pair u, v. The other two cells corner-adjacent to point p are empty.

Consider the grid graph $G = G_t(V_t, E_t)$, where the time-subscript will be omitted for convenience. Recall that G is a connected graph, so that $S(V)$, the union of the occupied cells (closed squares), is an arc-wise connected set in the plane. The complement $\overline{S}(V)$ of $S(V)$ consists of one or more connected components, called *holes*. They are denoted by $H_0, \ldots, H_k$, $k \geq 0$, where H_0 denotes the unique unbounded component, the *outer hole*, while every other hole H_i, $i \geq 1$, is said to be an *inner hole*. Each hole is bounded by a simple orthogonal polygon, which is called the *contour* of the hole.

The set of all cells in V side-adjacent or corner-adjacent to at least one cell in H_0, denoted B_{out}, is called the *outer boundary* of the configuration. Similarly, we define the *inner boundary* B_{in} of the configuration to be the set of all elements of V side or corner-adjacent to at least one empty cell belonging to an inner hole. Note that a cell may belong to both the outer boundary and the inner boundary of the configuration, and it may be adjacent to empty cells in more than one of the inner holes H_i, $i \geq 1$.

A hole H_i, $i \geq 0$, is said to be *critical* if it contains an empty cell side-adjacent to a pair of modules that form a critical pair. See Figure 3. A hole which is not critical is said to be *perfect*. It follows from the connectedness of our configurations that a critical pair is always associated with exactly two holes.

In what follows, we consider the set of *simple* (i.e., non-self-intersecting) cycles $\mathcal{C}$ in G. For any simple cycle $C \in \mathcal{C}$, let $R(C)$ denote the set of cells (empty or occupied) belonging to or enclosed by C. Define the *area* of C as $|R(C)|$. Each simple cycle C in the graph G corresponds to a simple closed curve obtained by connecting the centers of the cells belonging to C in the cyclic order. The area of C is the number of cells enclosed or crossed by this curve. The *contents* of a cycle $C \in \mathcal{C}$, denoted by $Q(C)$, is defined as $Q(C) = R(C) \cap V$. That is, $Q(C)$ is the set of occupied cells belonging to or enclosed by C. See Figure 3 for an illustration.

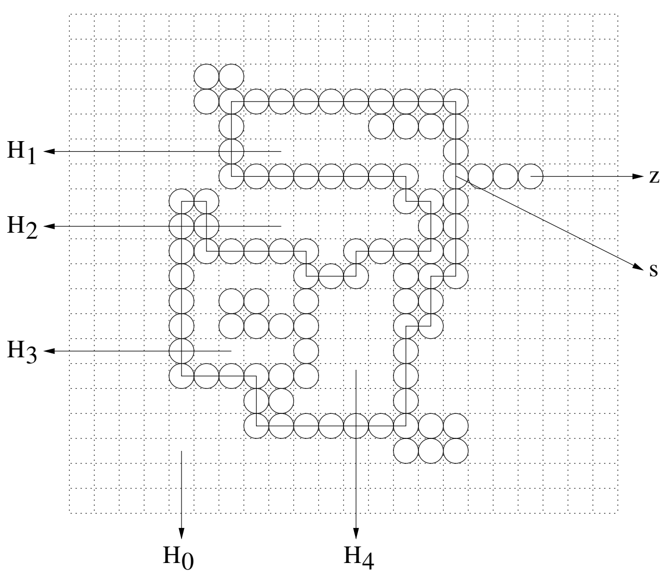

Figure 3: A configuration V in which the minimum degree in $V \setminus Z$ is at least two. There are three maximal cycles, having areas of 131, 4, and 6, and contents of size 89, 4, and 6. A polygonal line delineates the large cycle with area 131. There are five holes, out of which H_0 and H_2 are critical, and the other three are perfect.

We say that a cycle C is a *maximal cycle*, if $R(C)$ is maximal with respect to inclusion, i.e., there is no cycle C', such that $R(C) \subset R(C')$. Denote by $\mathcal{M}$ the set of maximal cycles in G. Note that if $\mathcal{C} \neq \emptyset$, $\mathcal{M} \neq \emptyset$. The next three lemmas give some useful properties of the system of maximal cycles, similar to those of the block decomposition of graphs [4].

Lemma 1. *If C and C' are two maximal cycles, then $|R(C) \cap R(C')| \leq 1$.*

Proof. Assume to the contrary that $|R(C) \cap R(C')| \geq 2$. Then $C \cup C'$ induces a (simple) cycle C'' in G so that $R(C) \subset R(C'')$ (and $R(C') \subset R(C'')$), contradicting the maximality of C (and C'). $\square$

Assume now that in the graph G the degree of every vertex except z is at least two at some point during the reconfiguration algorithm. Since G is connected, for any two maximal cycles, C and C', there exists a simple path $P = v_1 \ldots v_k$, $k \geq 1$, in G, connecting a vertex $v_1 \in C$ with $v_k \in C'$, none of whose intermediate vertices belong to C or C'. The vertex v_1 (resp. v_k) is said to be a *connector* for C (resp. for C'). (It is possible that v_1 and v_k coincide.) Notice that, although there may exist more than one such paths P connecting C and C', the connectors (their endpoints) are uniquely determined. See Figure 4. Similarly, for any maximal cycle $C \in \mathcal{M}$, consider a simple path connecting C to s (the base of Z), whose intermediate vertices do not belong to C, and let the corresponding vertex of C be also called a connector (that is, if $s \in C$, s is a connector). By the maximality of C, the connector is also unique in this case.

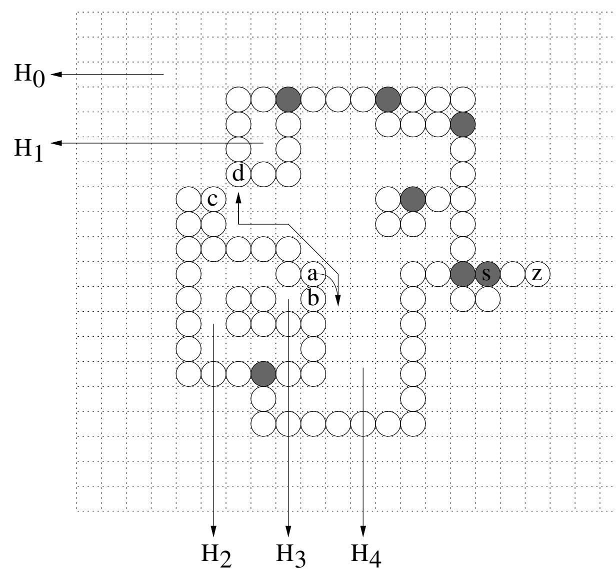

Figure 4: A configuration with five maximal cycles, with areas of 12, 8, 4, 4 and 39. Connectors are shaded in the figure. The configuration has five holes, four of which are critical (H_0, H_2, H_3, H_4). For example, c, d is a critical pair associated with holes H_0 and H_4. ab is a rightmost edge of the maximal cycle with area 39, as specified in Step 2 of the algorithm.

Lemma 2. *Assume that in G the degree of every vertex except of z is at least two. If $\mathcal{C} \neq \emptyset$, there exists a maximal cycle C having exactly one connector c. Furthermore, every vertex of G not in $R(C)$ can be connected to c by a path, all of whose intermediate vertices lie outside of $R(C)$.*

Proof. By definition, each maximal cycle has at least one connector. Assume for contradiction that each maximal cycle has at least two connectors. It follows from the degree condition that $\mathcal{C} \neq \emptyset$ and so $\mathcal{M} \neq \emptyset$. If $|\mathcal{M}| = 1$, i.e., $\mathcal{M} = C$, C has exactly one connector, as noted above. Assume therefore that $|\mathcal{M}| \geq 2$. There exist two maximal cycles, C_1 and C_2, and a simple path in G between them that does not pass through any vertex belonging to another maximal cycle. Since C_2 has at least two connectors, there exist two maximal cycles, C_2 and C_3, and a simple path between them that uses another connector of C_2 and does not include any point belonging to another maximal cycle. Any new cycle that we may reach, has at least two connectors, so we can repeat this procedure. Finally, we must either visit a vertex belonging to some path already visited before or reach one of the previously considered cycles. In either case, we obtain a simple cycle C' with $R(C_i) \subset R(C')$ for some $C_i \in \mathcal{M}$, which contradicts the maximality of C_i.

To verify the second part of the lemma, it is sufficient to observe that every vertex of G not in $R(C)$ lies either in $R(C')$ for some other element $C' \in \mathcal{M}$ or along a simple path connecting C to another maximal cycle $C' \in \mathcal{M}$, whose intermediate points do not lie in $R(C)$ or $R(C')$. $\square$

LEMMA 3. *Suppose that in G the degree of every vertex except of z is at least two, and let C be a maximal cycle with precisely one connector. Then no vertex of C other than its connector is adjacent in G to any other vertex that occupies a cell not in $R(C)$.*

PROOF. Suppose to the contrary that some vertex v_0 of C, different from its connector, is adjacent to another vertex $v_1 \in V \setminus Z$ that lies in the exterior of C. Let v_2 denote a neighbor of v_1 different from v_0. In general, if $v_i \neq s$ has already been determined for some $i \geq 2$, then let v_{i+1} be any neighbor of v_i in G, different from v_{i-1}. Using the fact that v_0 is not a connector, and using the maximality of C, we can argue that v_{i+1} cannot belong to C, and cannot be identical to s or to any other vertex v_j $(j \leq i)$. This procedure can be continued forever, which is impossible. $\square$

LEMMA 4. *Consider two empty cells, c and c', belonging to a hole H_i $(i \geq 0)$, each sharing at least one segment (side) with its contour. Place a new module m in cell c. Then there exists a sequence of moves through which m relocates to c' so that at each step m remains adjacent to the contour of H_i.*

PROOF. The assertion follows by analyzing all possible local configurations and by checking that the motion rules permit every single step of m making a full cycle along the contour of H_i. See Fig 5. $\square$

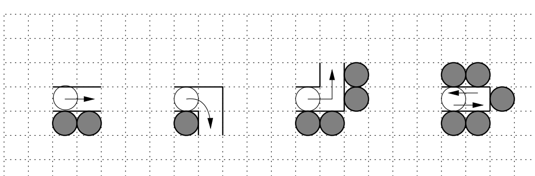

Figure 5: Moves along the contour of a hole. Stationary modules are shown shaded.

3.3 Algorithm description

Next, we describe one iteration of the algorithm, culminating in the extension of the chain Z by one module in the E-direction. For any hole H_i, let $\Delta(H_i)$ denote the set of all (empty) cells in H_i that contribute at least one segment to the contour of H_i.

STEP 1 Assume that there exists at least one vertex in $V \setminus Z$, whose degree in G is one, else go to STEP 2.

STEP 1.A: If there exists a vertex of degree one in $B_{\text{out}} \setminus Z$ (i.e., on the outer boundary), select one such module, say m. Clearly, the removal of m does not disconnect G, and once m is removed, $cell(m) \in \Delta(H_0)$, where H_0 is the new unbounded hole. By Lemma 4, if m is placed in $cell(m)$, there exists a sequence of moves through which m relocates to the empty cell $E(z)$, so that at each step $m \in \Delta(H_0)$. Thus, m extends Z by one cell. Then start a new iteration.

STEP 1.B: If no vertex of degree one in $V \setminus Z$ belongs to the outer boundary B_{out}, select a vertex m of degree one that belongs to the inner boundary B_{in}. The removal of m does not disconnect G, and once m is removed, $cell(m) \in \Delta(H_i)$ for some $i \geq 1$. By Lemma 4, if m is placed back in $cell(m)$, there is a sequence of moves taking m into a position where its degree in G is at least two. (The existence of such a position follows from the fact that H_i is an inner hole.) If now there exists a vertex of degree one in $B_{\text{out}} \setminus Z$, go to STEP 1.A. If there exists a vertex of degree one in B_{in}, repeat STEP 1.B.

STEP 2 Assume that in G the degree of every vertex belonging to $V \setminus Z$ is at least two. Then we have $\mathcal{C} \neq \emptyset$ so that $\mathcal{M} \neq \emptyset$. By Lemma 2, there exists a maximal cycle C having exactly one connector c. Consider all vertical edges of C with the smallest and with the largest x-coordinates. Assume without loss of generality that the connector of C does *not* belong to the highest vertical edge ab having the *largest* x-coordinate. Suppose further that a lies above b. (See Figure 6, where $\ldots u, a, b, v \ldots$ is the clockwise order of vertices of C in the neighborhood of ab. The case when the connector does not belong to the highest vertical edge of C having the *smallest* x-coordinate, can be treated similarly.) By Lemma 3, $E(a)$, $E(b)$, and $N(a)$ must be empty, and the removal of a does not disconnect the graph G. Set $m := a$. There are two cases, depending

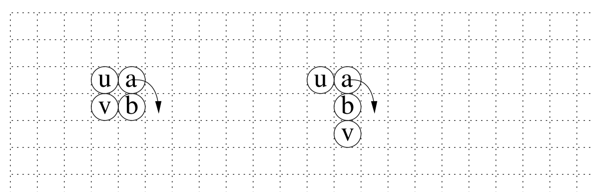

Figure 6: Illustration of possible moves in STEP 2 **of the algorithm.**

on whether m belongs to the outer boundary or not.

STEP 2.A: If $m \in B_{\text{out}} \setminus Z$, proceed as in STEP 1.A: relocate m to the empty cell $E(z)$ so that at each step $m \in \Delta(H_0)$. Thus m extends Z by one cell. Then start a new iteration.

STEP 2.B: If $m \notin B_{\text{out}}$, then we have $m \in B_{\text{in}}$, since m is adjacent to at least two empty cells, $E(a)$ and $N(a)$. Consider the holes H_i, H_j $(i, j \geq 1)$ including the E-side and the N-side of a, respectively. (It may happen that $i = j$.) It follows from the maximality of C that H_i and H_j cannot be perfect. Once m is removed, a becomes part of a larger hole $H = H_i \cup H_j \cup a$. Let $(c_1, d_1), (c_2, d_2), \ldots, (c_k, d_k)$ be the circular sequence of critical pairs around H, listed in clockwise order, and following a. That is, the cells $c_j, d_j \in V$ are corner-adjacent to each other and side-adjacent to an empty cell $e_j \in \Delta(H)$. See Fig 4. Obviously, $G \setminus \{m\}$ remains connected.

Claim. *There is exists an index j for which $\{c_j, d_j\} \not\subseteq C$, and so that in $G \setminus \{m\}$, there are two simple paths, $\pi(c_j)$ and $\pi(d_j)$, connecting c, the unique connector of C, to c_j and d_j, respectively, satisfying the following condition: $\pi(c_j)$, e_j, and $\pi(d_j)$, together with a piece of $C \setminus \{m\}$ induce a cycle C_j in G with $R(C_j) \supset R(C) \cup \{e_j\}$ and $|Q(C_j)| > |Q(C)|$.*

Let j denote the index whose existence is guaranteed by the Claim. In view of Lemma 4, if m is placed back in *cell*(m), then by a suitable sequence of moves it can be taken to cell e_j, where it is side-adjacent to both c_j and d_j. Therefore, in the final position, the area of C_j is larger than that of C and the contents of C_j is larger than that of C (since $\{c_j, d_j\} \not\subseteq C$). Then go to Step 1.

Proof of Claim. Assume without loss of generality that the cells $c_1, d_1, c_2, d_2, \ldots, c_k, d_k$ follow around the boundary of H in clockwise order, and set $T = \{c_1, d_1, c_2, d_2, \ldots, c_k, d_k\}$. Since the removal of m does not disconnect G, for any $x \in T$, there exists a simple path $\pi(x)$ in $G \setminus \{m\}$ connecting c to x. We can represent this path in the plane by a polygonal line through the centers of the corresponding cells. Let $\rho(x)$ denote a simple oriented Jordan arc running in the interior of the hole H from the *SW*-corner of a to the boundary of x. Connecting the endpoints of $\pi(x)$ and $\rho(x)$ by a simple arc in the union of the cells in $C \setminus \{a\}$ and by a segment in x, we obtain a closed Jordan region $J(x)$. Every cell $x \in T$ satisfies at least one of the following three possibilities:

1. x belongs to C,
2. $J(x)$ lies on the right-hand side of $\rho(x)$,
3. $J(x)$ lies on the left-hand side of $\rho(x)$.

It follows from the *maximality* of C that, for any $1 \le j \le k$, the cells d_j and c_{j+1} must be of the same type, where the indices are taken mod k. For similar reasons, as we go around the boundary of the hole H in clockwise direction, the first (resp. the last) cell that participates in a critical pair must be of type 1 or type 3 (resp. type 1 or type 2). By the definition of a, it cannot happen that both elements of every critical pair are of type 1. Thus, there exists a critical pair (c_j, d_j) whose elements have different types, and they will meet the requirements of the Claim. See Figure 7 for an illustration.

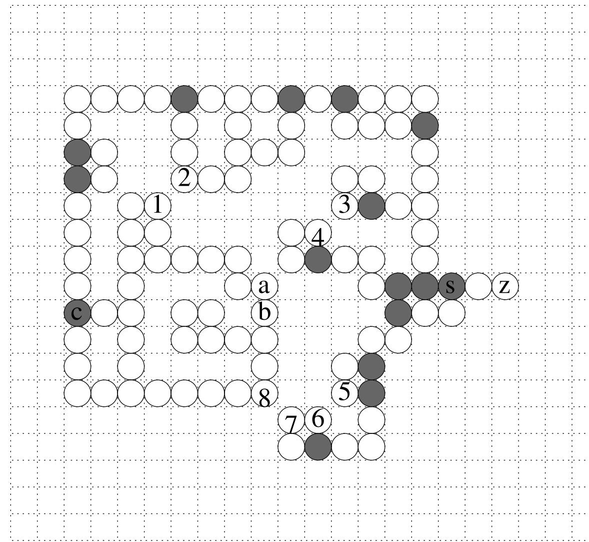

Figure 7: Illustration to the proof of Claim. A configuration with nine maximal cycles, with areas of 47, 18, 8, 6, 4, 4, 4, 4 and 4. Connectors are shaded in the figure. ab is a rightmost edge of the maximal cycle with area 47, as specified in Step 2 of the algorithm. Its unique connector is c. The critical pairs of the hole H created by the removal of a are: $(1, 2)$, $(3, 4)$, $(5, 6)$ and $(7, 8)$. Cells 1 through 7 have type 3, while cell 8 has types 1 and 2. Critical pair $(7, 8)$ meets the requirements of the Claim.

3.4 Algorithm analysis

Each of the steps 1.A, 1.B, 2.A, or 2.B consists of at most n moves by a single module. We prove that after performing at most n steps of type 1.B and at most n steps of type 2.B, we must make a step of type 1.A or 2.A, which completes one iteration of the algorithm.

First, note that the sequence of relocations in Step 1.B (if not empty) strictly increases the total number of edges in the graph G. As G remains connected throughout the algorithm, its number of edges is always at least $n - 1$. On the other hand, G is a subgraph of the infinite grid with $n \ge 2$ vertices, so its number of edges never exceeds $2n - 3$. Taking into account that the number of edges does not decrease during Step 2.B, we can conclude that each iteration uses fewer than n steps of type 1.B.

Recall that the contents of a cycle C, denoted $Q(C)$, is the set of occupied cells belonging to or enclosed by C. Consider the following weight function characterizing the configuration: $\Phi(G) = q - y$, where $q = |\cup_{C \in \mathcal{M}} Q(C)|$ is the total contents of maximal cycles, and y is the number of maximal cycles. Observe that $0 \le \Phi(G) \le n$. Note that Step 1.B preserves all the existent cycles, and observe that the above weight function strictly increases during each step of type 2.B, and it does not decrease during any step of type 1.B. Consider first a step of type 1.B: y can only increase by 1, and if y increases by 1, then q also increases by at least one (by the maximality condition). Consider now a step of type 2.B: Since $\{c_j, d_j\} \not\subseteq C$, assume for simplicity that $d_j \notin C$. Denote by q' and y' the values of q and y respectively at the end of the step. If $d_j \in Q$, we have $q' \ge q$ and $y' < y$; whereas if $d_j \notin Q$, we have $q' > q$ and $y' \le y$. In either case, $\Phi(G)$ strictly increases. Therefore, each iteration uses at most n steps of type 2.B.

In conclusion, after fewer than $n + n = 2n$ steps of type 1.B or 2.B, we must perform at least one step of type 1.A or 2.A, and complete an iteration of the algorithm. After each iteration, the horizontal chain Z is extended by one cell in the E-direction, so after $n-1$ iterations, the reconfiguration to a straight chain is complete. Each step consists of at most

n moves, so the entire algorithm takes fewer than $2n^3$ moves.

4. ANOTHER RECTANGULAR MODEL

The following two generic motion rules (Figure 8) define the *weak rectangular model*. These are to be understood as possible in all axis-parallel orientations, in fact generating eight rules, four diagonal moves and four side moves (axis-parallel ones). The only imposed condition is that the configuration must remain connected at each discrete time step.

- *Diagonal move*: A module m moves diagonally to an empty cell corner-adjacent to $cell(m)$.
- *Side move*: A module m moves to an empty cell side-adjacent to $cell(m)$.

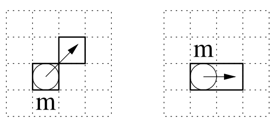

(a) (b)

Figure 8: (a) NE diagonal move and (b) side move in the E direction. The cells in which the moves take place are outlined in the figure.

The same result as in Theorem 1 holds for this second model, but its proof is much easier.

THEOREM 2. *The set of motion rules of the weak rectangular model guarantees the feasibility of motion planning for any pair of connected configurations having the same number of modules.*

PROOF. Assume without loss of generality that the lowest cell of the configuration has $y = 0$. Consider the following weight function characterizing the configuration:

$$\Phi(G) = \sum_{c \in V} |y(c)|,$$

where $y(c)$ is the y-coordinate of (the center of) cell c. Notice that Φ is invariant with respect to horizontal translation. If $\Phi(G) = 0$, then G must be a straight horizontal chain. We show that there exists a sequence of moves during which $\Phi(G)$ monotone decreases to 0, with the additional condition that at each time step the y-coordinate of every vertex is nonnegative.

Consider the top row of V; if its y-coordinate is equal to 0, there is nothing to prove (the reconfiguration is complete). Else consider its rightmost cell u. Notice that $N(u)$, $NW(u)$ and $NE(u)$, as well as $E(u)$, are all empty. Since G is connected, at least one of the cells $W(u)$ or $S(u)$ is nonempty. We have seven cases, five for the first alternative and two for the second. See Figure 9.

For the first five cases, $v := W(u)$ is nonempty.

Case 1: $S(v)$ is nonempty and $S(u)$ is empty. Then u makes a side move in the S direction.

Case 2: $S(v)$ and $S(u)$ are nonempty, and $SE(u)$ is empty. Then u makes a side move in the SE direction.

Case 3: $S(v)$, $S(u)$ and $SE(u)$ are nonempty. Then u makes a side move in the E direction.

Case 4: $S(v)$ is empty and $S(u)$ is nonempty. Then u makes a diagonal move in the SW direction.

Case 5: both $S(u)$ and $S(v)$ are empty. Then u makes a diagonal move in the SW direction.

For the last two cases, $W(u)$ is empty and $w := S(u)$ is nonempty.

Case 6: $SE(u)$ is empty. Then u makes a diagonal move in the SE direction.

Case 7: $SE(u)$ is nonempty. Then u makes a side move in the E direction.

It is easy to check that the configuration remains connected after each move. The reader should notice that the key is case 4 (a move which is prohibited in the previous model). Note that each of the cases 1, 2, 4, 5, and 6 reduces the weight function by one unit. Cases 3 or 7 can occur in a sequence at most n times after which one of the other five cases (1, 2, 4, 5, or 6) must occur. Therefore, $\Phi(G) = 0$ after at most $n^3/2$ steps (the initial value of $\Phi(G)$ is at most $\sum_{i=1}^{n-1} i \le n^2/2$).

A more careful calculation shows that in fact not more than $O(n^2)$ moves are made. Among all modules with minimum y-coordinate, the one whose x-coordinate is the smallest is called the *reference module* of the configuration. Consider the reference module, say r, of the initial configuration. Observe that the reconfiguration procedure outlined above does not move any of the modules in the lowest row, so in particular r remains fixed. In the only cases when $y(u)$ is not reduced (remains constant), 3 and 7, $x(u)$ is increased by one. Note also that $x(u)$ is decreased by one only if $y(u)$ is decreased by one (cases 4 and 5). By the connectedness of each intermediate configuration, $x(r) - n \le x(m) \le x(r) + n$, for each module m. Consequently, the number of moves of the module from cell c of the initial configuration that leave the value of y unchanged, is not more than $2n + y(c)$. Since the number of moves of the same module, during which y gets reduced is $y(c)$, the total number of moves in the reconfiguration process does not exceed

$$\sum_{c \in V} (2n + y(c) + y(c)) = 2n^2 + 2\sum_{c \in V} y(c) \le 3n^2.$$

This bound is tight up to a multiplicative constant, see Section 5. It is very easy to improve the last bound to $2n^2$. $\square$

As for the previous model, it is not hard to construct examples that cannot be reconfigured to straight chains using only diagonal moves or only side moves. There are, in fact, small examples that simultaneously work for both models.

5. CONCLUDING REMARKS

1. In our original model, the reconfiguration of a vertical chain into a horizontal one requires only $\Theta(n^2)$ moves, and we believe that no other pair of configurations requires more. We have shown that this holds in the weak rectangular model.

2. A somewhat different model can be obtained if, instead of the connectedness requirement at each time step, one imposes the following so-called *single backbone* condition [6]. Denote by $F_t = V_t \setminus \{m\}$ the set of modules that do *not* move in step t (i.e., fixed). We refer to F_t as the *backbone*. The initial configuration V_0 must be connected: i.e., the graph $G_0 = (V_0, E_0)$

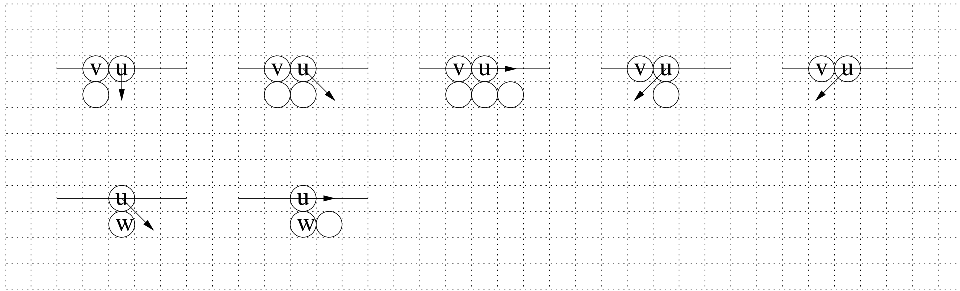

Figure 9: Case analysis in the proof of Theorem 2. In cases 1 through 5 (upper row), the cell $W(u)$ is nonempty; in cases 6 and 7 (bottom row), the cell $W(u)$ is empty. The horizontal line indicates in each case the top row of the configuration.

must be connected. At any time t, the backbone F_t must be connected with respect to side-adjacency, i.e., the graph $B_t = (F_t, E'_t)$ must be connected, where $E'_t \subset E_t$ is the set of edges connecting pairs of cells in F_t that are side-adjacent. (The backbone could however have 'holes'.) Together with the connectedness of the modules at time 0, it ensures that the modules remain connected at any time step t. (If concurrent moves are allowed, additional conditions have to be imposed, as in [6].)

A subtle difference exists between requiring the configuration to be connected at each discrete time step and requiring the existence of a connected backbone along which a module slides or rotates [6]. A one step motion that does not satisfy the single backbone condition appears in Figure 10: the initial connected configuration practically disconnects during the move and reconnects at the end of it.

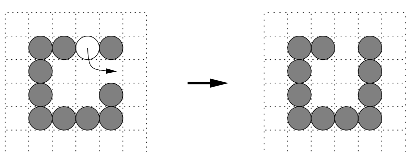

Figure 10: A rotation move which temporarily disconnects the configuration.

Notice that at each step of our algorithm, the element m selected to move has the property that its removal does not disconnect G. This property is also maintained throughout the sequence of moves until another element is selected to move. Therefore, the single backbone condition remains satisfied during the whole procedure.

3. A related question is the following. A configuration consisting of unit cubes of integer coordinates in d-space is called an *animal* if the boundary of their union is homeomorphic to a $(d-1)$-sphere. It is easy to see that in the plane, any animal can be transformed into any other by adding or removing one square at a time so that all intermediate configurations are animals. The corresponding statement in higher dimensions is not known to be true. There exist, however, relatively small animals in 3-space with the property that no cube can be removed from them without violating the condition. This is in sharp contrast to the situation in the plane.

6. REFERENCES

[1] Z. Butler, K. Kotay, D. Rus and K. Tomita, Generic decentralized control for a class of self-reconfigurable robots, *Proceedings of the 2002 IEEE International Conference on Robotics and Automation* (ICRA '02), Washington, May 2002, 809–816.

[2] G. Chirikjian, A. Pamecha and I. Ebert-Uphoff, Evaluating efficiency of self-reconfiguration in a class of modular robots, *Journal of Robotic Systems*, **13(5)** (1996), 317–338.

[3] E. Demaine, M. Demaine and H. Verrill, Coin-moving puzzles, in *More Games of No Chance*, edited by R. J. Nowakowski, pp. 405-431, Cambridge University Press, 2002.

[4] R. Diestel, *Graph theory (2nd ed.), Graduate Texts in Mathematics* **173**, Springer-Verlag, New York, 2000.

[5] A. Dumitrescu, I. Suzuki and M. Yamashita, Formations for fast locomotion of metamorphic robotic systems, *International Journal of Robotics Research*, to appear. A preliminary version in *Proceedings of the 2002 IEEE International Conference on Robotics and Automation* (ICRA '02), Washington, May 2002, 123–128.

[6] A. Dumitrescu, I. Suzuki and M. Yamashita, Motion planning for metamorphic systems: feasibility, decidability and distributed reconfiguration, *IEEE Transactions on Robotics and Automation*, to appear.

[7] S. Murata, H. Kurokawa and S. Kokaji, Self-assembling machine, *Proceedings of IEEE International Conference on Robotics and Automation*, (1994), 441–448.

[8] A. Nguyen, L. J. Guibas and M. Yim, Controlled module density helps reconfiguration planning, *Proceedings of IEEE International Workshop on Algorithmic Foundations of Robotics*, 2000.

[9] A. Pamecha, I. Ebert-Uphoff and G. Chirikjian, Useful metrics for modular robot motion planning, *IEEE Transactions on Robotics and Automation*, **13(4)** (1997), 531–545.

[10] D. Rus and M. Vona, Crystalline robots: self-reconfiguration with compressible unit modules, *Autonomous Robots*, **10** (2001), 107–124.

[11] M. Yim, Y. Zhang, J. Lamping and E. Mao, Distributed control for 3D metamorphosis, *Autonomous Robots*, **10** (2001), 41–56.

[12] E. Yoshida, S. Murata, A. Kamimura, K. Tomita, H. Kurokawa and S. Kokaji, A motion planning method for a self-reconfigurable modular robot, in *Proceedings of the 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2001.