Unfolding Manhattan Towers [1]

Mirela Damian [a] and Robin Flatland [b] and Joseph O'Rourke [c,2]

[a] *Dept. of Computer Science, Villanova University, Villanova, PA 19085, USA.*

[b] *Dept. of Computer Science, Siena College, Loudonville, NY 12211, USA.*

[c] *Dept. of Computer Science, Smith College, Northampton, MA 01063, USA.*

Abstract

We provide an algorithm for unfolding the surface of any orthogonal polyhedron that falls into a particular shape class we call Manhattan Towers, to a nonoverlapping planar orthogonal polygon. The algorithm cuts along edges of a $4 \times 5 \times 1$ refinement of the vertex grid.

Key words: Unfolding, orthogonal, genus-zero, polyhedron.

1 Introduction

It is a long-standing open problem to decide whether the surface of every convex polyhedron can be *edge unfolded*: cut along edges and unfolded flat to one piece without overlap [DO05][DO07]. It is known that some nonconvex polyhedra have no edge unfolding; a simple example is a small box sitting on top of a larger box. However, no example is known of a nonconvex polyhedron that cannot be unfolded with unrestricted cuts, i.e., cuts that may cross the interior of faces.

The difficulty of these questions led to the exploration of *orthogonal polyhedra*, those whose faces meet at right angles and whose edges are parallel to coordinate axes. Progress has been made in two directions: firstly, by restricting the shapes to subclasses of orthogonal polyhedra, such as the "orthostacks"

Email addresses: mirela.damian@villanova.edu (Mirela Damian), flatland@siena.edu (Robin Flatland), orourke@cs.smith.edu (Joseph O'Rourke).

[1] A preliminary version of this paper appeared in CCCG 2005.

[2] Supported by NSF Distinguished Teaching Scholars award DUE-0123154.

and "orthotubes" studied in [BDD+98]; and secondly, by generalizing the cuts beyond edges but with some restrictions. In particular, a *grid unfolding* partitions the surface of the polyhedron by coordinate planes through every vertex, and then restricts cuts to the resulting grid. The box-on-box example mentioned earlier can be easily grid unfolded. Recent work on grid unfolding of orthostacks is reported in [DM04] and [DIL04].

Because on the one hand no example is known of an orthogonal polyhedron that cannot be grid unfolded, and on the other hand no algorithm is known for grid unfolding other than very specialized shapes, the suggestion was made in [DO04] to seek $k_1 \times k_2 \times k_3$-*refined grid unfoldings*, where every face of the vertex grid is further refined into a grid of edges. Positive integers k_1, k_2 and k_3 are associated with the amount of refinement in the x, y and z directions, respectively; e.g., z-perpendicular faces are refined into a $k_1 \times k_2$ grid, and similarly x-perpendicular (y-perpendicular) faces are refined into a $k_2 \times k_3$ ($k_1 \times k_3$) grid. It is this line we pursue in this paper, on a class of shapes not previously considered.

We define "Manhattan Tower (MT) polyhedra" to be the natural generalization of "Manhattan Skyline polygons," also known as rectilinear histogram polygons [O'R87, p. 176]. Although we do not know of an unrefined grid unfolding for this class of shapes, we prove (Theorem 2) that there is a $4 \times 5 \times 1$ grid unfolding. Our algorithm peels off a spiral strip that winds first forward and then interleaves backward around vertical slices of the polyhedron, recursing as attached slices are encountered.

2 Definitions

Let Z_k be the plane $\{z = k\}$, for $k \geq 0$. Define $\mathcal{P}$ to be a *Manhattan Tower* (MT) if it is an orthogonal polyhedron such that:

(1) $\mathcal{P}$ lies in the halfspace $z \geq 0$, and its intersection with Z_0 is a simply connected orthogonal polygon, its base polygon;
(2) For $0 \leq k < j$, $\mathcal{P} \cap Z_k \supseteq \mathcal{P} \cap Z_j$: the cross-section at higher levels is nested in that for lower levels.

Manhattan Towers are *terrains* in that they meet each vertical (parallel to z) line in a single segment or not at all; thus they are *monotone* with respect to z. Fig. 1a shows an example. Manhattan Towers may not be monotone with respect to x or y, and indeed $\mathcal{P} \cap Z_k$ will in general have several connected components (see Fig. 2c), and may have holes (see Fig. 2b), for $k > 0$.

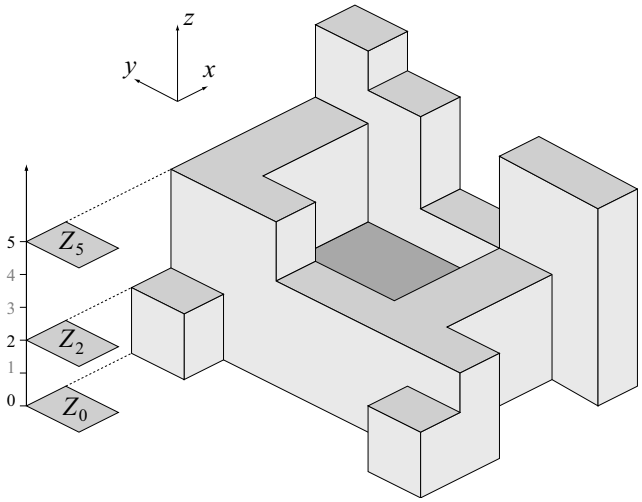

Fig. 1. Manhattan Tower $\mathcal{P}$.

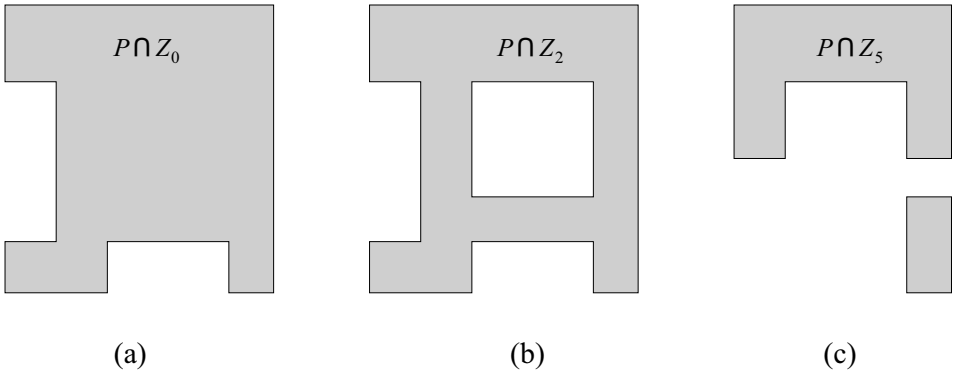

Fig. 2. Cross-sections of Manhattan Tower P from Fig. 1: (a) The base $Z_0 \cap P$ is a simple orthogonal polygon; (b) $Z_2 \cap P$ is an orthogonal polygon with one hole; (c) $Z_5 \cap P$ has two disjoint components.

As an xy-plane sweeps from Z_0 upwards, the cross-section of $\mathcal{P}$ changes at finitely many locations. Thus a Manhattan Tower $\mathcal{P}$ may be viewed as consisting of nested layers, with each layer the extrusion of a set of orthogonal polygons. The *base layer* of $\mathcal{P}$ is its bottom layer, which is bounded below by Z_0 and above by the xy-plane passing through the first vertex with $z > 0$. Note that, unlike higher layers, the base is simply connected, since it is an extrusion of $\mathcal{P} \cap Z_0$.

We use the following notation to describe the six types of faces, depending on the direction in which the outward normal points: *front*: $-y$; *back*: $+y$; *left*: $-x$; *right*: $+x$; *bottom*: $-z$; *top*: $+z$. An *x-edge* is an edge parallel to the x-axis; *y-edges* and *z-edges* are defined similarly.

Clockwise (cw) and counterclockwise (ccw) directions are defined with respect to the view from $y = -\infty$. Later we will rotate the coordinate axes in recursive calls, with all terms tied to the axes altering appropriately.

3 Recursion Tree

We start with the partition Π of the base layer induced by the xz-planes passing through every vertex of $\mathcal{P}$. (The restriction of the partition to planes orthogonal to y will facilitate processing in the $\pm y$ directions.) Such a partition consists of rectangular boxes only (see Fig. 3a). The dual graph of Π has a node for each box and an edge between each pair of nodes corresponding to adjacent boxes. Since the base is simply connected, the dual graph of Π is a tree T (Fig. 3b), which we refer to as the *recursion tree*. The root of T is a node corresponding to a box (the *root box*) whose front face has a minimum y-coordinate (with ties arbitrarily broken).

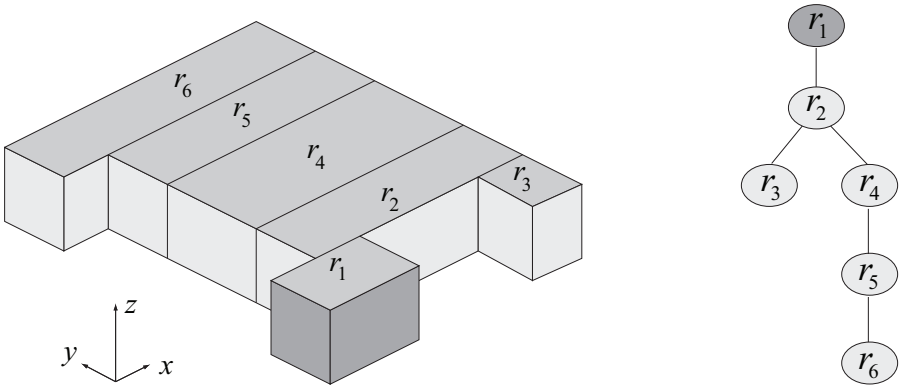

(a) (b)

Fig. 3. (a) Partition Π of $\mathcal{P}$'s base; (b) Recursion tree T.

It turns out that nearly all unfolding issues are present in unfolding single-layer MTs, due to the nested-layer structure of MTs. In Sec. 5 we describe an algorithm for unfolding single-layer MTs. The algorithm is then extended to handle multiple-layer MTs in Sec. 6.

4 $4 \times 5 \times 1-$Refined Manhattan Towers

Fig. 4 illustrates the refinement process, using the base from Fig. 3a as an example. The *gridded* base (Fig. 4a) contains additional surface edges induced by yz-coordinate planes through each vertex. A $4 \times 5 \times 1$ refinement of the gridded base further partitions each horizontal grid rectangle into a 4×5 grid. In addition to gridedges of the gridded base, the $4 \times 5 \times 1$-*refined* base (Fig. 4b) contains all surface edges induced by coordinate planes passing through each gridpoint in the refinement. In the following we show that every $4 \times 5 \times 1$-refined MT can be edge-unfolded.

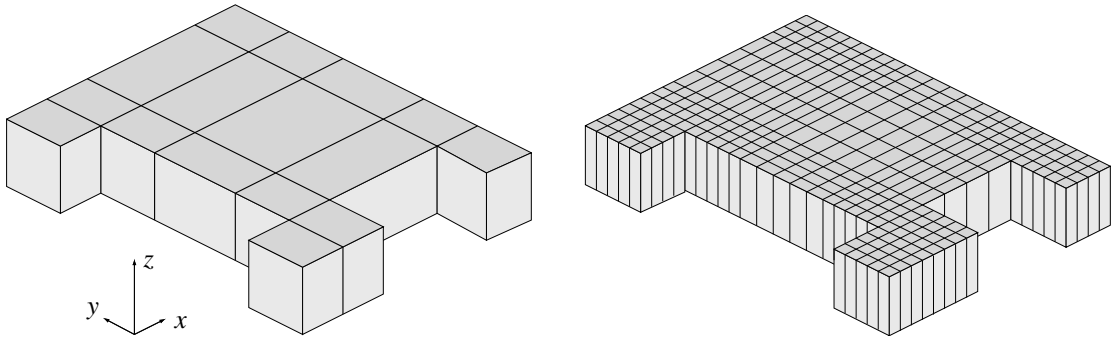

(a) (b)

Fig. 4. (a) Gridded MT base; (b) $4 \times 5 \times 1$-refined MT base.

5 Single-Layer MTs

A single-layer Manhattan Tower consists of a single layer, the *base* layer. We describe the unfolding algorithm recursively, starting with the base case in which the layer is a single rectangular box.

5.1 Single Box Unfolding

Let r be a $4\times5\times1$-gridded rectangular box and let T, R, B, L, K and F be the top, right, bottom, left, back and front faces of r, respectively. Let s and t be two gridpoints either adjacent along an x-edge of r (as in Fig. 5a), or vertically aligned, with one on a top x-edge and one on a bottom x-edge of the front face of r (as in Fig 6a). Let y_s and y_t be the (y parallel) gridedges incident to s and t. The unfolding of r starts at y_s and ends at y_t. More precisely, this means the following. Let ξ_{2d} (ξ_{3d}) denote the planar (three-dimensional) embedding of the cut surface piece. Then ξ_{2d} has y_s on its far left and y_t on its far right (as in Figs. 5c and 6c).

The main unfolding idea is to cut the top, right, bottom and left faces so that they unfold into a staircase-like strip, and then attach front and back faces to it vertically without overlap. We collectively refer to the top, right, bottom and left faces as *support* faces (intuitively, they support the front and back faces). Roughly stated, ξ_{3d} starts at y_s, spirals cw around the support faces toward the back face, crosses the back face, then spirals ccw around the support faces back to y_t. This idea is illustrated in Figs. 5 and 6. In the following we provide the details for the case when s and t are adjacent on the top front edge of r (Fig. 5). The case when s lies on a bottom front edge and t lies on a top front edge of r is similar and is illustrated in Fig. 6; the case when s is on the top and t is on the bottom is identical, when viewed through an xy-mirror.

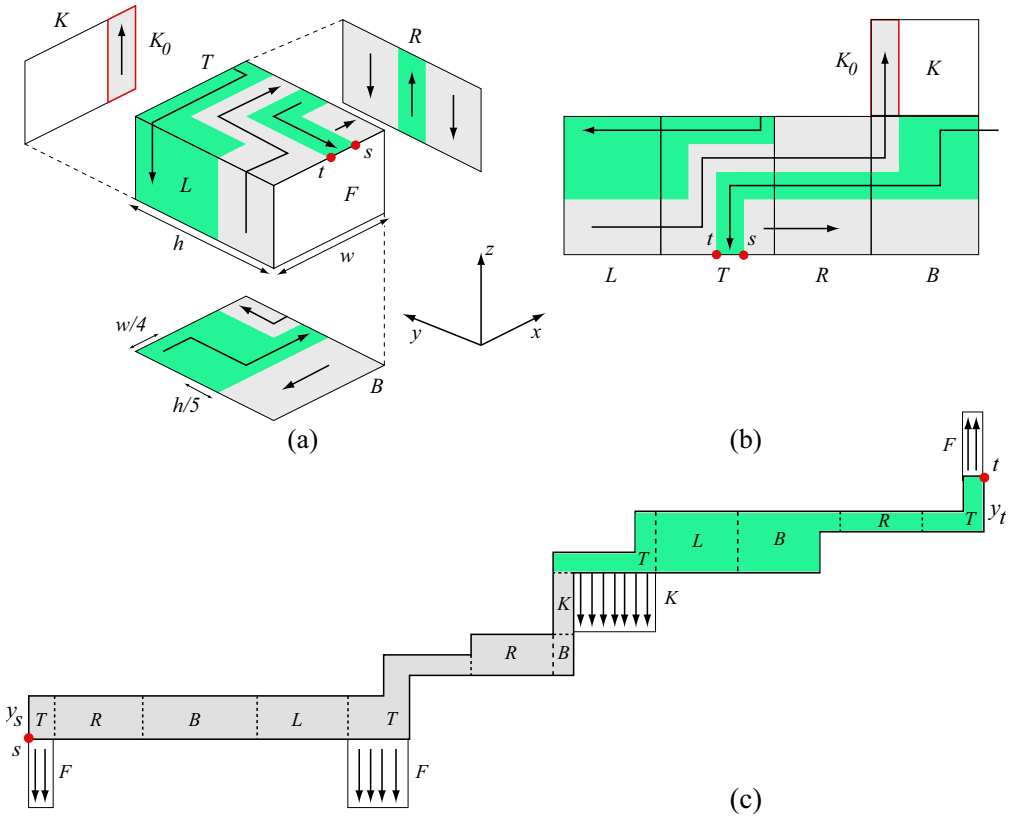

Fig. 5. Single box unfolding: s adjacent to t along an x-edge (a) Front view of box r and mirror view of right (R), bottom (B) and back (K) faces, marked with unfolding cuts (b) Faces of r flattened out (front face not shown) (c) Spiral unfolding of r; labels identify faces containing the unfolded pieces.

As illustrated in Fig. 5, let w be the x-extent and let h be the y-extent of r. We implicitly define the unfolding cuts by describing the surface pieces encountered in a walk along ξ_{3d} on the surface of r (delineated by unfolding cuts). Starting at y_s, walk cw along a rectangular strip of y-extent equal to $2h/5$ (two gridfaces wide) that spirals around the support faces from y_s to y_t. This spiral strip lies adjacent to the front face of r; we will refer to it as the *front spiral* of ξ_{3d}. At y_t, take a left turn and continue along a rectangular strip (orthogonal to the front spiral and right-aligned at t) of y-extent equal to $2h/5$ (two gridfaces wide) and x-extent equal to $w/4$ (one gridface long). At the end of this strip, take a right turn and continue along a rectangular strip of y-extent equal to $h/5$, until the right face R is met; at this point, the strip thickens to a y-extent equal to $2h/5$ (two gridfaces wide), so that it touches the back face K of the box. The strip touching K consumes the entire length of the right face R, plus an additional $w/4$ (one gridface) amount onto the adjacent bottom face B. At the end of this bottom strip, take a left turn and continue along a $w/4$-wide strip across back face K and up onto the top face T. The piece of ξ_{3d} traversed so far is called the *forward* spiral; the remaining piece is called the *backward* spiral, conveying the fact that from this

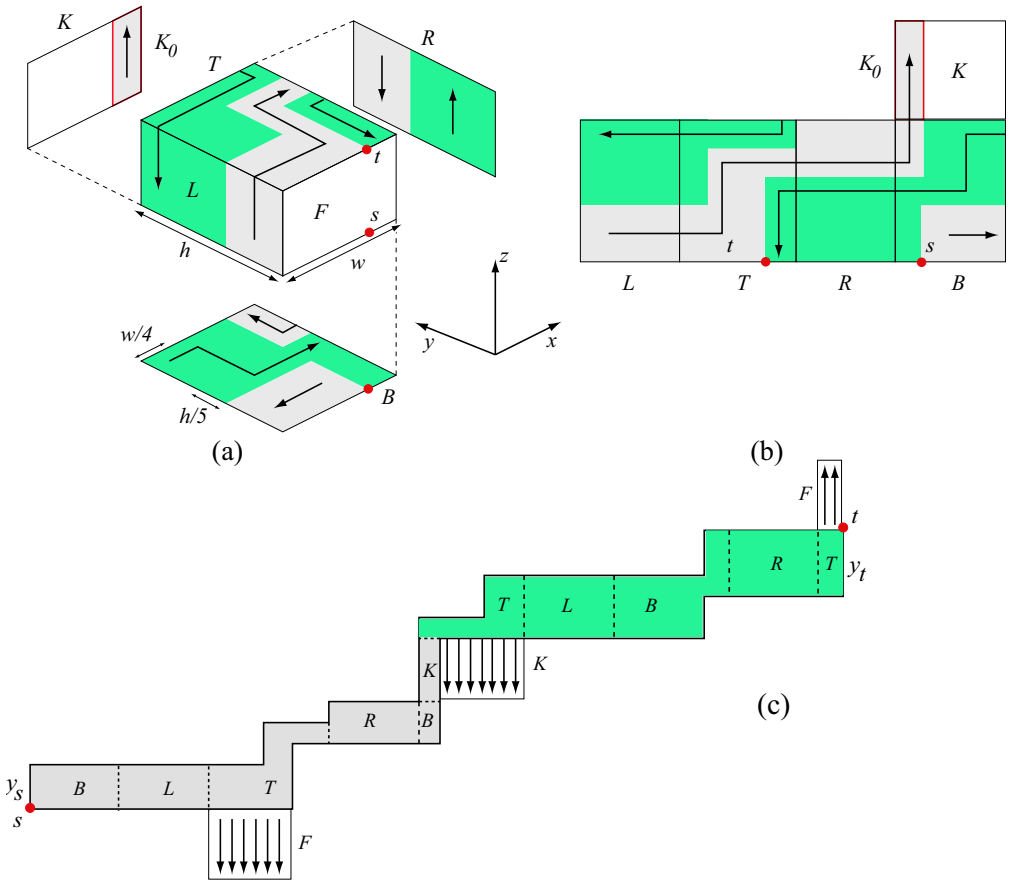

Fig. 6. Single box unfolding: s and t are vertically aligned on opposite front edges (a) Front view of box r and mirror view of right (R), bottom (B) and back (K) faces, marked with unfolding cuts (b) Faces of r flattened out (front face not shown) (c) Spiral unfolding of r; labels identify faces containing the unfolded pieces.

point on ξ_{3d} spirals ccw around the support faces back to y_t. The piece of the backward spiral adjacent to the back face is the *back spiral* of r. The planar piece ξ_{2d} (obtained by laying ξ_{3d} out in the plane) has the staircase-like shape illustrated in Fig 5c. Conceptually, the front face F and the back face K are not part of the unfolding described so far; however, they can be flipped up and attached vertically to ξ_{2d} without overlap (see the arrowed faces in Fig. 5c), a point to which we return in Sec. 5.4.

5.2 Recursion Structure

In general, a box r in the recursion tree has children (adjacent boxes) attached along its front and/or back face. Call a child attached on the front a *front child* and a child attached on the back a *back child*. In unfolding r, we unwind the support (top, bottom, left, right) faces into a staircase-like strip just as described for the single box. But when the front/back spiral runs alongside

the front/back face of r and encounters an adjacent child, the unfolding of r is temporarily suspended, the child is recursively unfolded, then the unfolding of r resumes where it left off.

At any time in the recursive algorithm there is a *forward* direction, corresponding to the initial spiraling from front to back (the lighter strip in Figs. 5 and 6), and an opposing *backward* direction corresponding to the subsequent reverse spiraling from back to front (the darker strip in Figs. 5 and 6). When the recursion processes a front child, the sense of forward/backward is reversed: we view the coordinate system rotated a half-turn about the vertical axis so that the $+y$ axis is aligned with the forward direction of the child's spiral, with all terms tied to the axes altering appropriately. In particular, this means that the start and end unfolding points s', t' of a front child r' lie on the front face of r', as defined in the rotated system.

For example, in Fig. 7, boxes a, b, c, d are processed from front to back. But recursion on e, a front child of d, reverses the sense of forward, which continues through e, f, and g. We can view the coordinate system rotated so that $+y$ is aligned with the arrows shown. Thus f is a back child of e, g is a back child of f, and k a front child of g. Again the sense of forward is reversed for the processing of k.

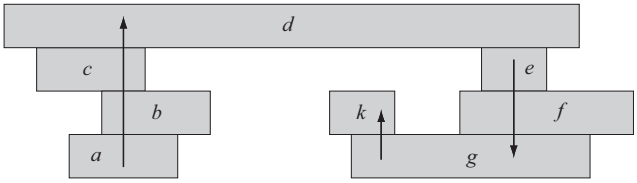

Fig. 7. Arrows indicate which direction is *forward* in the recursive processing.

5.3 Suturing Techniques

We employ two methods to "suture" a child's unfolding to its parent's unfolding. The first method, *same-direction suture*, is used to suture all front children to their parent. As the name suggests, this suturing technique preserves the unwinding direction (cw or ccw) of the parent's spiral. If there are no back children, then a strip from the back face of the parent (K_0 in Figs. 5 and 6) is used to reverse the direction of the spiral to complete the parent's unfolding, as described in Sec. 5.1 for the single box. However, if the parent has one or more back children, these children cover parts or perhaps all of the back face of the parent, and the back face strip may not be available for the reversal. So instead we use a second suturing method, *reverse-direction suture*, for one of the back children. This suture uses the child's unfolding to reverse the direction of the parent's spiral, and does not require a back-face strip. We choose

exactly one back child for reverse-direction suturing. Although any such child would serve, for definiteness we select the rightmost child. To summarize, our suturing rules are as follows:

(1) For every front child, use same-direction suturing.
(2) For the rightmost back child, use reverse-direction suturing.
(3) For remaining back children, use same-direction suturing.

5.3.1 Same-direction suture

We first note that a front child r' never entirely covers the front face of its parent box r, because the parent of r is also adjacent to the front face of r. This is evident in Fig. 7, where e cannot cover the front face of d because d's parent, c, is also adjacent along that side. Similarly, k cannot cover the "front" face of g (where here the sense of front is reversed with the forward direction of processing) because g's parent f is also adjacent along that side. The same-direction suture may only be applied in such a situation of non-complete coverage of the shared front face, for it uses a thin (one-gridface-wide) vertical strip of that face.

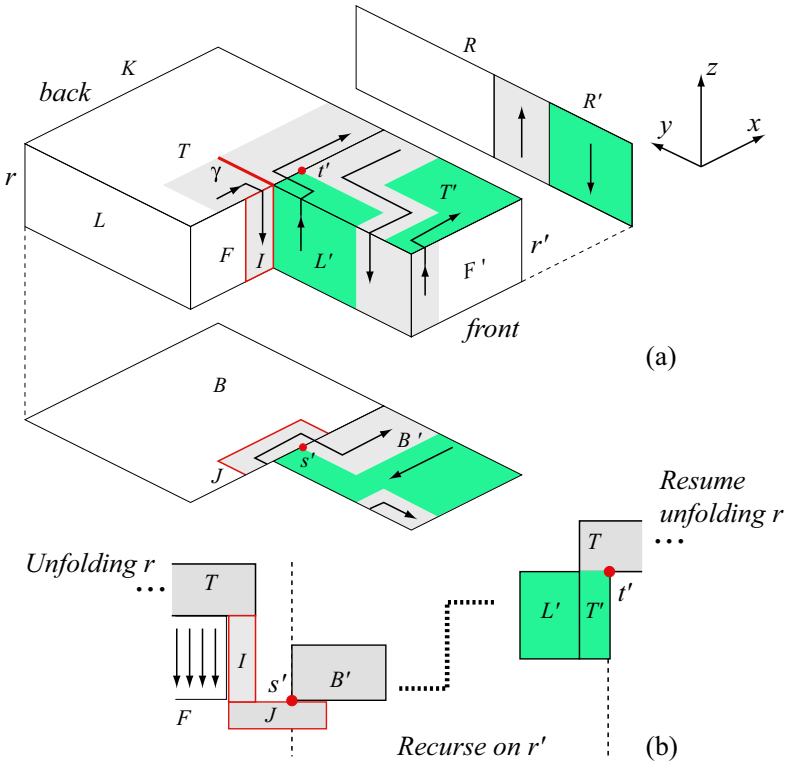

Fig. 8. Same-direction suture. (a) Front view of root box r and front child r', with mirror bottom, left and back views. (b) Result ξ_{2d} of recursive unfolding.

This suture begins at the point where the parent's spiral meets an adjacent child as it runs alongside its front or back face. To be more specific, consider

the case when r' is a front child of r, and the parent's front spiral meets r' as it runs along the top of r. This situation is illustrated in Fig. 8. The same-direction suture begins by cutting a vertical strip I off the front face of parent r, which includes all vertical gridfaces that lie alongside child r' (see Fig. 8a), then it takes a bite J one gridface thick and three gridfaces long (in the x-direction) off the bottom face of the parent. This marks the gridedge $y_{s'}$ on r' where the child's spiral unfolding starts. The child's spiral unfolding ends at top gridedge $y_{t'}$ of the same x-coordinate as $y_{s'}$. When the child's unfolding is complete, the spiral unfolding of the parent resumes at the y-gridedge it left off (see the cut labeled γ in Fig. 8). The other cases are similar: if r' is a back child of r, I occurs on the back face of r; and if the parent's front spiral meets r' as it runs along the bottom of r,[3] J occurs on the top face of r (see child r_4 and parent r_2 in Fig. 10). It is this last case that requires a 5 refinement in the y direction: the front spiral must be two gridfaces thick so that cutting J off it does not disconnect it.

In Fig. 8, notice that the parent's spiral unfolds in cw direction on top face T before the suture begins. The parent's cw unfolding is suspended at y-gridedge marked γ, and after the child is unfolded, the parent's spiral resumes its unfolding in cw direction at γ. The unfolded surface ξ_{2d} is shown in Fig. 8b.

5.3.2 Reverse-direction suture

This suture begins after the parent's spiral completes its first cycle around the support (top, right, bottom, left) faces, as illustrated in Fig. 9 for parent r and back child r'. As in the single box case (Sec. 5.1), after a forward move in the $+y$-direction, the spiral starts a second cycle around the support faces. However, unlike in the single box case, the spiral stops as soon as it reaches a y-gridedge of the same x-coordinate as the rightmost gridpoint u that the parent shares with a back child. At that point, the parent's spiral continues with a gridface-thick strip S in the $+y$-direction, right-aligned at y_u. Let s' be the left corner of S on the boundary shared by r and r'. The unfolding of r' begins at gridedge $y_{s'}$ and ends at gridedge $y_{t'}$ immediately to the left of $y_{s'}$ on top of r'. When the child's unfolding is complete, the unfolding of the parent resumes at the gridedge it left off, with the spiral unwinding in reverse direction.

As the name suggests, this suturing technique reverses the unwinding direction (cw or ccw) of the parent's spiral. In Fig. 9, notice that the parent's spiral unfolds in cw direction on top face T before the suture begins. After the child is unfolded, the parent's spiral resumes its unfolding in ccw direction at $y_{s'}$. The result ξ_{2d} of this unfolding is shown in Fig. 9b.

[3] This only happens if r' is a front child of r.

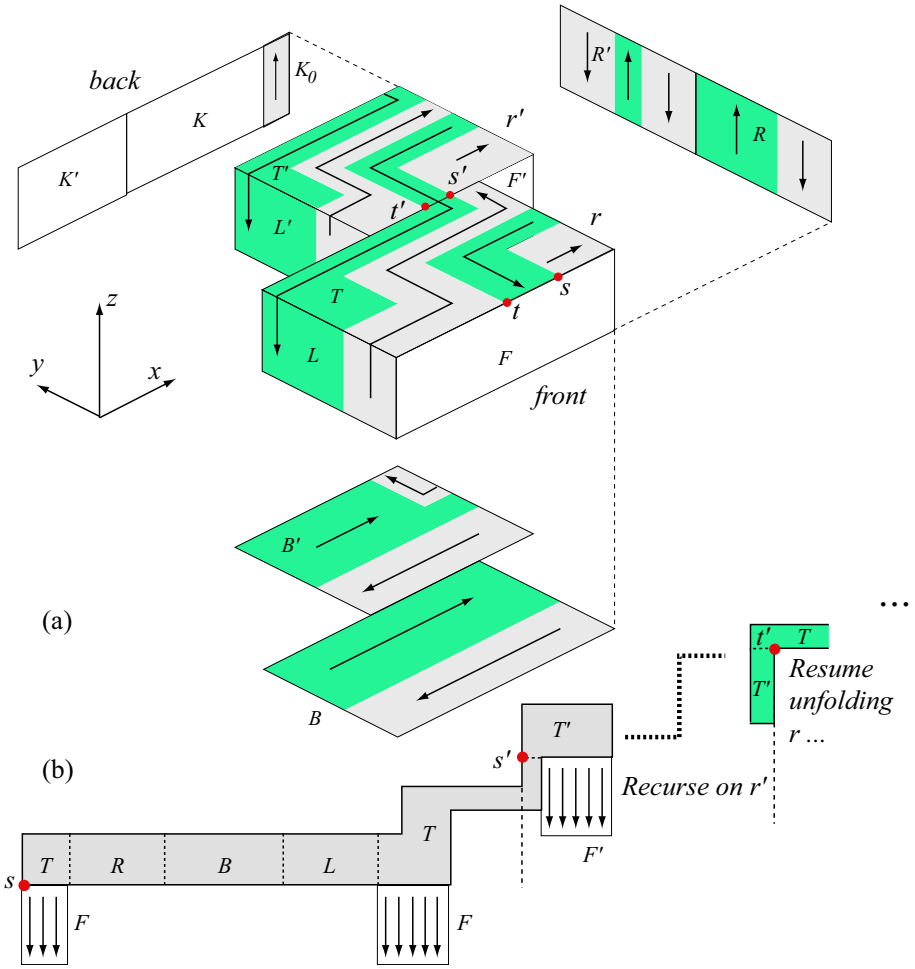

Fig. 9. Reverse-direction suture. (a) Front view of faces of root box r and back child r', with mirror bottom, left and back views. (b) Result ξ_{2d} of recursive unfolding.

5.4 Attaching Front and Back Faces

The spiral strip ξ_{3d} covers all of the top, bottom, right, and left faces of the base. It also covers the gridface-thick strips of a front/back face used by the same-direction sutures (I in Fig. 8) and the gridface-thick strips of back faces used to reverse the spiral direction in the base cases (K_0 in Figs. 5 and 6). The staircase structure of ξ_{2d} (shown formally in Theorem 1) guarantees that no overlap occurs.

We now show that remaining exposed front and back pieces that are not part of ξ_{3d} can be attached orthogonally to ξ_{2d} without overlap. Consider the set of top gridedges shared by top faces with front/back faces. These gridedges occur on the horizontal boundaries of ξ_{2d} as a collection of one or more contiguous segments. We partition the front/back faces by imagining these top gridedges emanate downward lightrays on front/back faces. Then all front and back pieces are illuminated, and these pieces are attached to correspond-

ing illuminating gridedges (see Figs. 5c, 6c, 8b and 9b). Although no interior points overlap in the unfolding, we allow *edge overlap*, which corresponds to the physical model of cutting out the unfolded piece from a sheet of paper. For example, in Fig. 9b a left gridedge of F' overlaps a gridedge of ξ_{2d}. It is not difficult to avoid edge overlap (e.g. by making the portion of the strip causing the edge overlap narrower to separate it from F'), but doing so requires increasing the degree of refinement.

The next section summarizes the entire unfolding process for single-layer MTs.

5.5 Unfolding Algorithm for Single-Layer MTs

Consider an arbitrary base partitioned into rectangular boxes with xz-planes $Y_0, Y_1, \ldots$ through each vertex. Select a root box r adjacent to Y_0 (breaking ties arbitrarily) and set the forward unwinding direction d to be cw. Let y_s and y_t be top y-gridedges of r, as described in Sec. 5.1 for the single-box case. Our recursive unfolding starts at root box r and proceeds as follows.

`Algorithm UNFOLD`(r, y_s, y_t)

---

1. Start unfolding the forward spiral piece adjacent to front face (§ 5.1).

2. **Unfolding Front Children.** For each front child r' of r encountered
   Determine gridedges $y_{s'}$, $y_{t'}$ using same-direction suture (§ 5.3.1).
   Recurse: UNFOLD$(r', y_{s'}, y_{t'})$.

3. If r has no back children then complete the unfolding of r (§ 5.1) and exit.
4. Determine start and end gridedges $y_{s'}$, $y_{t'}$ for rightmost back child r' using reverse-direction suture (§ 5.3.2).
5. Complete the unfolding of the forward spiral up to $y_{s'}$ (§ 5.3.2).
6. Recurse: UNFOLD$(r', y_{s'}, y_{t'})$.
7. Continue unfolding the back spiral adjacent to back face (§ 5.1).

8. **Unfolding Rest of Back Children.** For each back child r' of r encountered
   Determine gridedges $y_{s'}$, $y_{t'}$ using same-direction suture (§ 5.3.1).
   Recurse: UNFOLD$(r', y_{s'}, y_{t'})$.

9. Complete the unfolding of r by spiraling back to y_t (§ 5.1).
10. Hang front and back faces off the unfolded spiral. (§ 5.4).

This algorithm can be easily implemented to run in $O(n^2)$ time on a polyhedron $\mathcal{P}$ with n vertices. Fig. 10 illustrates the recursive unfolding algorithm

on a 3-legged H-shaped base. The unfolding starts at gridedge y_{s_1} of root box r_1 and ends at gridedge y_{t_1}. (Only the endpoints s_1 and t_1 of these two gridedges are marked in Fig 10.) The spiral strip encounters the boxes in the order $r_1, r_2, r_3, r_4, r_5, r_6$ and r_7, which corresponds to the ordering of the recursive calls. For each i, y_{s_i} and y_{t_i} are gridedges of r_i where the unfolding of r_i starts and ends. The algorithm uses reverse-direction suture to attach back child r_2 to parent r_1; same-direction suture to attach front child r_3, and then r_4, to parent r_2; reverse-direction suture to attach back child r_5 to parent r_2; and same-direction suture to attach back child r_6, and then r_7, to parent r_2. Note that a refinement of 5 in the y direction is necessary on top of box r_2 for this unfolding.

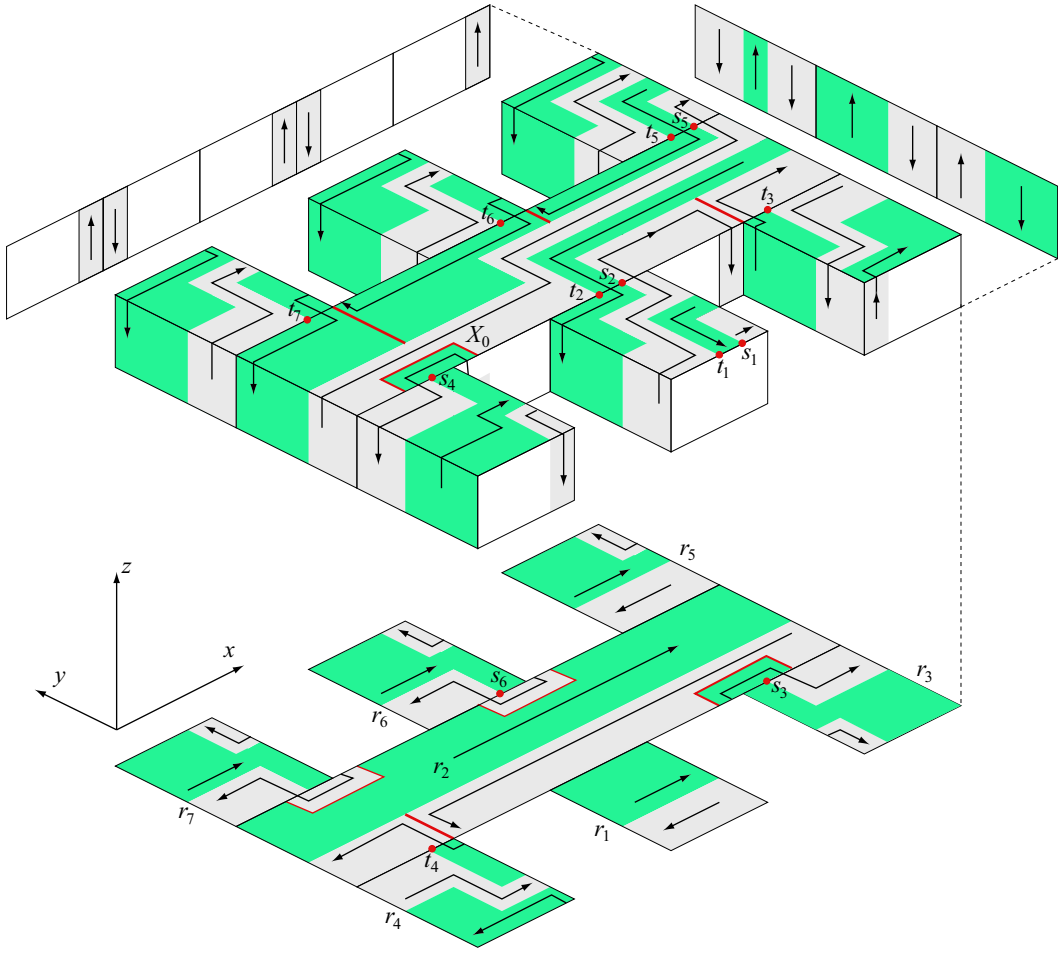

Fig. 10. Unfolding a 3-legged H-shaped base.

Theorem 1 *The* `UNFOLD`(r, y_s, y_t) *algorithm unfolds all boxes in the recursion tree rooted at* r *into a staircase-like strip* ξ_{2d} *completely contained between the vertical lines passing through* y_s *and* y_t.

PROOF. The proof is by induction on the height k of the recursion tree rooted at r. The base case is $k = 0$ and corresponds to single node trees. This is the case illustrated in Figs. 5 and 6, which satisfy the claim of the theorem.

The inductive hypothesis is that the theorem is true for any recursion tree of height $k-1$ or less. To prove the inductive step, consider a recursion tree T of height k rooted at r. The staircase strip $\xi_{2d}(r)$ of r alone, ignoring all children, fits between the vertical lines passing through y_s and y_t (see Figs. 5c and 6c).

Assume without loss of generality that r unfolds cw. There are two possible placements of s and t on r: (i) s and t are on *opposite* top/bottom edges of the front face of r (Fig. 6a), as placed by a same-direction suture, or (ii) s and t are on a *same* top/bottom edge of r (Fig. 5a), as placed by a reverse-direction suture. In either case, s and t are placed in such a way that no children exist along the path extending cw from t to s on r. This means that all front children of r are encountered during the unwinding of r's front spiral from s to t on r. That all back children are encountered during the unwinding of r's back spiral is clear: starting at the rightmost back child, the back spiral makes a complete cycle around the back face.

Consider now an arbitrary child r' of r in T and let T' be the subtree rooted at r'. As noted above, r' will be encountered during the unfolding of r. Let $y_{s'}$ and $y_{t'}$ be the gridedges on r' where the unfolding of r' starts and ends. The inductive hypothesis applied on T' tells us that the strip $\xi_{2d}(r')$ corresponding to T' fits between the vertical lines passing through $y_{s'}$ and $y_{t'}$. Fig. 8b illustrates the same-direction suture: when $\xi_{2d}(r')$ is sutured to $\xi_{2d}(r)$, the strip $\xi_{2d}(r)$ expands horizontally and remains contained between the vertical lines passing through y_s and y_t. The reverse-direction suture has a similar behavior (illustrated in Fig. 9b), thus completing this proof.

6 Multiple-Layer MTs

Few changes are necessary to make the single-layer unfolding algorithm from Sec. 5.5 handle multiple-layer Manhattan Towers. In fact, the view of the cuts used to form ξ_{3d} from $z = \pm\infty$ in the multi-layer case is identical to that in the single-layer unfolding. All the differences lie in vertical (z-parallel) strips used to adjust for differing tower heights (here we use the term "tower" to refer to a rectangular prism sitting on a box of the gridded MT base). When there are multiple layers, the basic unit to unfold is a vertical *slab* $S(r)$ consisting of a box r in the partition Π of the base layer and all the towers that rest on top of r (see Fig. 11). A slab is a Manhattan Skyline polygon parallel to the xz-plane extruded in the y direction: the projection of the top faces of the slab on the xy-plane forms a partition of the (unique) bottom face (face

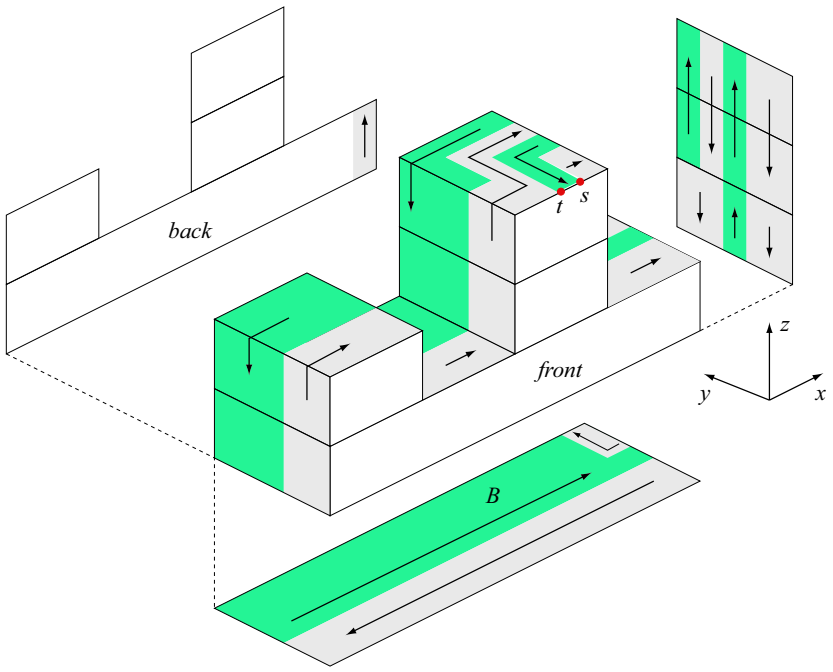

Fig. 11. Front view of single slab $S(r)$, with mirror bottom, left and back views.

B in Fig. 11). It is here that we make essential use of the assumptions that $\mathcal{P} \cap Z_0$ is a simply connected orthogonal polygon, and the cross-sections at higher levels are nested in those for lower levels.

The unfolding of a slab $S(r)$ is similar to the unfolding of a single box:

(1) Select an arbitrary top face T of the slab.
(2) Select start and end gridedges y_s and y_t on T as in the single box case.
(3) Unfold $S(r)$ using the procedure described in Sec. 5.1 for r.

The only difference is that a slab may have multiple left/right/top faces, causing the spiral ξ_{3d} to cycle up and down over the towers of $S(r)$. Fig. 11 illustrates this for the case when both y_s and y_t lie on the top front edge of $S(r)$. As a result, ξ_{2d} lengthens horizontally, but still maintaining its staircase structure. As in the case of a single box, ξ_{3d} covers all of the top, right, bottom and left faces. The remaining front and back pieces are attached to ξ_{2d} using the illumination scheme described in Sec. 5.4. In general, a multiple-layer MT $\mathcal{P}$ consists of many slabs; in this case, we use the recursion tree for the base of $\mathcal{P}$ to unfold $\mathcal{P}$ recursively (in this sense, single-layer and multiple-layer MTs have identical recursion structures). The recursive unfolding algorithm is similar to the algorithm described in Sec. 5.5 for single-layer MTs, with some minor modifications to accommodate the existence of towers. In the following we describe these modifications with the help of the MT example from Fig. 12, whose base is the 3-legged H-shape single-layer MT from Fig. 10.

Let $S(r')$ be the slab corresponding to a child r' of r. When the unfolding

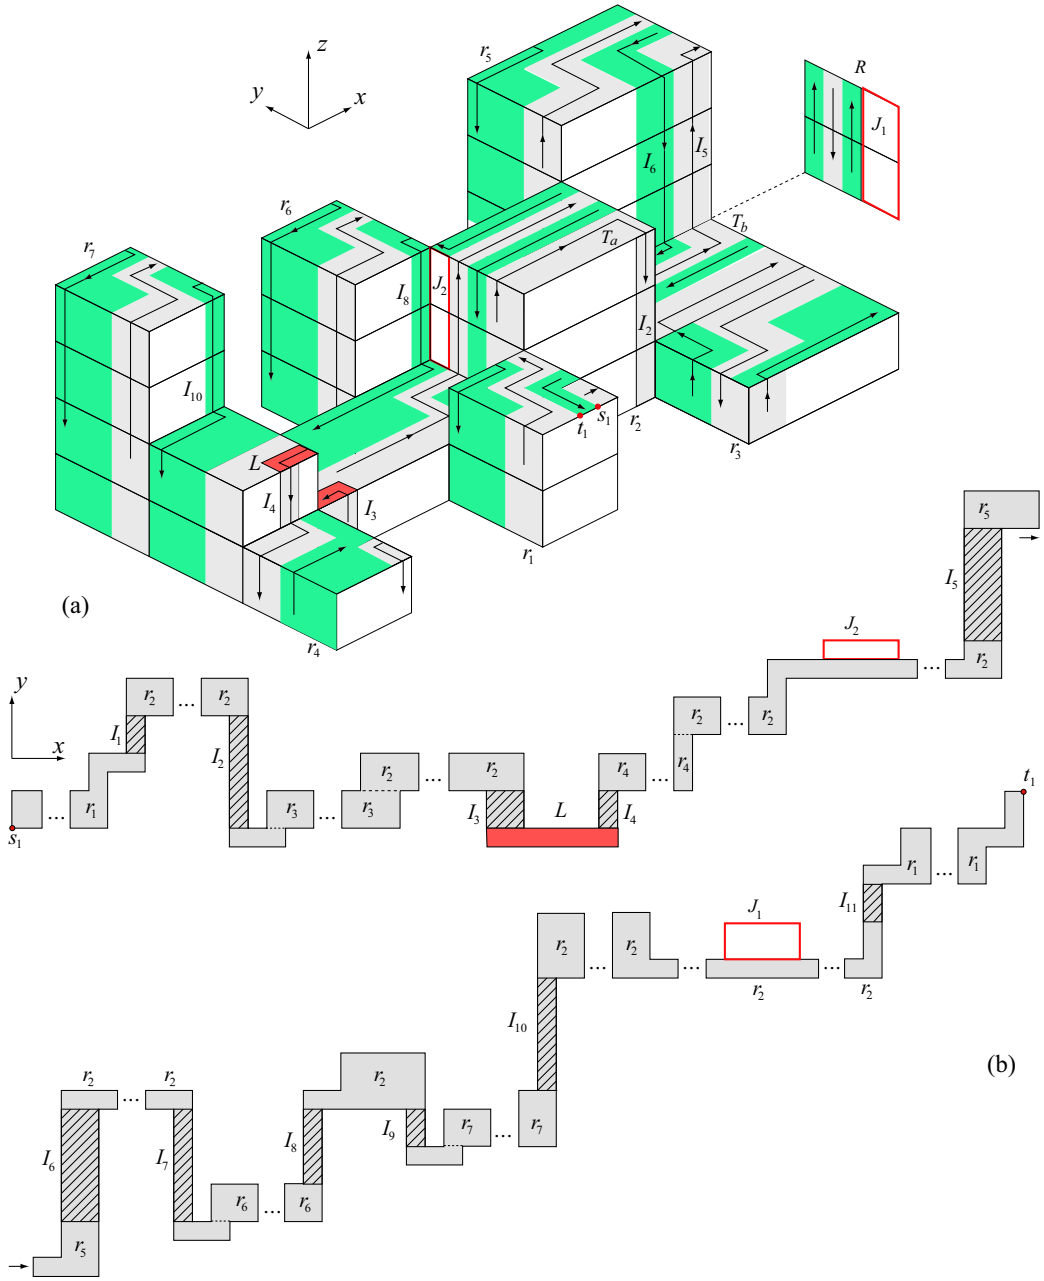

Fig. 12. Unfolding multiple-layer MTs. (a) Spiral ξ_{3d}; bottom and back mirror views are as shown in Fig. 10 (b) ξ_{2d}, strips J_1 and J_2 attached above; transitions between towers are striped; piece labels correspond to MT boxes to which they belong.

strip for $S(r)$ first encounters a top/bottom face f of $S(r')$ (when viewed from $z = +\infty$), the unfolding of $S(r)$ is suspended in favor of $S(r')$. Next we discuss the two suturing techniques used to glue the unfolding of $S(r')$ to the unfolding of $S(r)$.

Same-direction suture. In this case, the bottom/top face opposite to f is used to accommodate the start unfolding gridedge $y_{s'}$ for $S(r')$; the end

unfolding gridedge $y_{t'}$ is selected as in the single-layer case.

Consider first the case when r' is a front child of r. If $S(r')$ is encountered while ξ_{3d} runs along the top of $S(r)$, the suture is identical to the single-layer case: a vertical strip across the front of $S(r)$ is used to reach the bottom of $S(r')$ (see strip I_2 in Fig. 12, reaching front child $S(r_3)$). If $S(r')$ is encountered while ξ_{3d} runs along the bottom of $S(r)$, the suture is similar to the single-layer case, with two simple modifications:

(1) After using a vertical strip to reach the top of $S(r)$, a small "bite" is taken out of the top of $S(r)$ to reach the top of $S(r')$ in the single-layer case. In the multiple-layer case, it may be necessary to extend such a bite up/down a z-face in order to reach the point of the same x-coordinate as $y_{s'}$. This is the case of slab $S(r_4)$ in Fig. 12: strip I_3 is used to get from the bottom of $S(r_2)$ to the top of $S(r_2)$, after which the "bite" labeled L extends up a right face of $S(r_2)$ to reach the x-coordinate of y_{s_4}.
(2) Unlike the single-layer case, a top bite used in the same-direction suture is not necessarily adjacent to child $S(r')$. In this case, a second z-strip (such as I_4 in Fig. 12) is used to reach the top of $S(r')$.

The case in which r' is a back child of r is similar and is illustrated in Fig. 12: strips I_7 and I_9 (visible in Fig. 12b, but not in 12a) are used to make the transition from $S(r_2)$ to $S(r_6)$ and $S(r_7)$ respectively, and strips I_8 and I_{10} are used to return to $S(r_2)$.

Reverse-direction suture. As in the same-direction suture case, a vertical strip may be needed to make transitions between the top of a parent $S(r)$ and the top of a child $S(r')$ that uses reverse-direction suture. This is the case for $S(r_3)$ in Fig. 12, where the vertical strip I_6 (I_7) is used to move from (to) $S(r_2)$ to (from) $S(r_5)$.

The result of these alterations is that ξ_{2d} may lengthen vertically, but it remains monotone in the horizontal direction.

One final modification is necessary due to the difference in height between towers that belong to a same slab (see for instance towers T_a and T_b of $S(r_2)$ in Fig. 12a). In such cases it is possible that the spiral ξ_{3d} does not completely cover the left/right faces of the slab. We resolve this problem by thickening ξ_{3d} in the y-direction to cover the uncovered pieces. To be more precise, consider the vertical strip labeled J_1 in Fig. 12 (in the mirror view of right face R). The reason J_1 remains uncovered is because in unfolding $S(r_3)$, the unfolding of $S(r_2)$ suspends at the top y-gridedge of J_1 and resumes at the bottom y-gridedge of J_1. Similarly, ξ_{3d} skips over the strip marked J_2 in Fig. 12: when the back spiral of $S(r_2)$ meets $S(r_6)$, the unfolding of $S(r_2)$ suspends at the top y-gridedge of J_2 and resumes at the bottom y-gridedge of J_2.

We resolve the problem of uncovered strips as follows. First, note that every uncovered strip is on a left/right face (never a back/front face) of a slab. This means that each left/right piece of ξ_{3d} adjacent to an uncovered strip can be thickened until it completely covers it. This results in vertically thicker pieces in the planar embedding ξ_{2d} of ξ_{3d}. Because ξ_{2d} is monotonic in the horizontal direction, thickening it vertically cannot result in overlap. It also cannot interfere with the hanging of the front/back faces, since front/back faces attach along horizontal (x-parallel) sections of ξ_{3d}, whereas the thickened strips occur along otherwise unused vertical (z-parallel) sections of ξ_{3d}. Thus we have the following result.

Theorem 2 *Every Manhattan Tower polyhedron can be edge-unfolded with a* $4 \times 5 \times 1$ *refinement of each face of the vertex grid.*

7 Conclusion

We have established that every $4 \times 5 \times 1$-refined Manhattan Tower polyhedron may be edge-unfolded. This is the second nontrivial class of objects known to have a refined grid-unfolding, besides orthostacks. This is the first unfolding algorithm for orthogonal polyhedra that uses recursion, something we believe will be useful in developing algorithms to unfold more general shapes that can branch in many directions. The algorithm works on some orthogonal polyhedra that are not Manhattan Towers, and we are exploring how to widen its range of applicability. Finally, we note that if the Manhattan tower base polygon is a rectangle (rather than an arbitrary orthogonal polygon), then a nonrecursive $1 \times 1 \times 1$ (i.e., unrefined) grid-unfolding algorithm is recently available [O'R07].

Acknowledgements. We thank the anonymous referees for their careful reading and insightful comments.

References

[BDD+98] T. Biedl, E. D. Demaine, M. L. Demaine, A. Lubiw, J. O'Rourke, M. Overmars, S. Robbins, and S. Whitesides. Unfolding some classes of orthogonal polyhedra. In *Proc. 10th Canad. Conf. Comput. Geom.*, pages 70–71, 1998.

[DIL04] E. D. Demaine, J. Iacono, and S. Langerman. Grid vertex-unfolding of orthostacks. In *Proc. Japan Conf. Discrete Comp. Geom.*, November 2004. To appear in LNCS, 2005.

[DM04] M. Damian and H. Meijer. Grid edge-unfolding orthostacks with orthogonally convex slabs. In *14th Annu. Fall Workshop Comput. Geom.*, pages 20–21, November 2004.

[DO04] E. D. Demaine and J. O'Rourke. Open problems from CCCG 2004. In *Proc. 16th Canad. Conf. Comput. Geom.*, 2004.

[DO05] E. D. Demaine and J. O'Rourke. A survey of folding and unfolding in computational geometry. In J. E. Goodman, J. Pach, and E. Welzl, editors, *Combinatorial and Computational Geometry*. Cambridge University Press, 2005.

[DO07] E. D. Demaine and J. O'Rourke. *Geometric Folding Algorithms: Linkages, Origami, Polyhedra*. Cambridge University Press, 2007. `http://www.gfalop.org`.

[O'R87] J. O'Rourke. *Art Gallery Theorems and Algorithms*. The International Series of Monographs on Computer Science. Oxford University Press, New York, NY, 1987.

[O'R07] J. O'Rourke. Unfolding orthogonal terrains. Technical Report 084, Smith College, July 2007. arXiv:0707.0610v4 [cs.CG].