Partitioning Orthogonal Histograms into Rectangular Boxes

Therese Biedl[1], Martin Derka[2], Veronika Irvine[1], Anna Lubiw[1], Debajyoti Mondal[3], and Alexi Turcotte[1]

[1] Cheriton School of Computer Science, Univ. of Waterloo, Waterloo, Canada
[2] School of Computer Science, Carleton University, Ottawa, Canada
[3] Department of Computer Science, University of Saskatchewan, Saskatoon, Canada

Abstract. The problem of partitioning an orthogonal polyhedron into a minimum number of boxes was shown to be NP-hard in 1991, but no approximability result is known except for a 4-approximation algorithm for 3D-histograms. In this paper we broaden the understanding of the 3D-histogram partitioning problem. We prove that partitioning a 3D-histogram into a minimum number of boxes is NP-hard, even for histograms of height two. This settles an open question posed by Floderus et al. We then show the problem to be APX-hard for histograms of height four. On the positive side, we give polynomial-time algorithms to compute optimal or approximate box partitions for some restricted but interesting classes of polyhedra and 3D-histograms.

1 Introduction

Partitioning a geometric object or a shape into simpler parts is a classic problem in computational geometry. Such partitioning problems are motivated by their applications in image processing, camera placement in security systems, computer graphics, VLSI manufacturing, and so on.

An important special case is when the object is an *orthogonal* polygon or polyhedron—meaning that the edges or faces are parallel to the axes or coordinate planes—and the goal is to partition into a minimum number of rectangles or *boxes*, where a *box* is an orthogonal polyhedron with 6 faces (i.e., the 3D equivalent of a rectangle).

In two dimensions the problem of partitioning an orthogonal polygon into a minimum number of rectangles can be solved in polynomial time, both for simple polygons and for polygons with holes [14, 13, 12, 7, 5]. In three dimensions the problem becomes NP-hard in general, as proved by Dielissen and Kaldewaij [3]. Our aim is to explore the boundary between hard and easy for a special class of orthogonal polyhedra, namely *histograms*.

A *2D-histogram* is an orthogonal polygon L that contains an edge e such that for any point $p \in L$, the line segment connecting p to its orthogonal projection on e lies entirely inside L. See Fig. 1(a). Similarly, a *3D-histogram* is an orthogonal polyhedron H that contains a face f such that for any point $p \in H$, the line segment connecting p to its orthogonal projection on f lies entirely inside H.

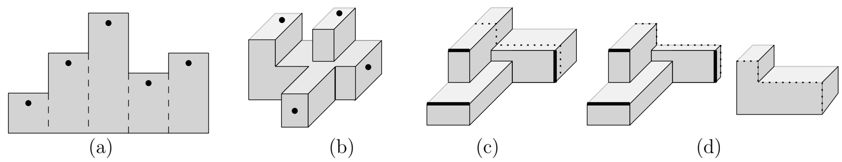

Fig. 1. (a) A 2D-histogram with a minimum partition into rectangles, as witnessed by the given “independent” points, no two of which lie in a rectangle. (b) A 3D-histogram requiring 4 boxes in its optimal partition since there are four “independent” points (no two in a box). (c)–(d) A guillotine cut.

See Fig. 1(b). The face f is called the *base* of the histogram. Note that any histogram can be adjusted so that the vertices (and consequently also edges and faces) have integer coordinates, and the combinatorial structure of the histogram is preserved. Throughout this paper, we assume that any histogram has integer coordinates and that the base has z-coordinate 0 and all other faces lie above the base. The *height* of a face parallel to the base is its z-coordinate, and the maximum of these is called the *height* of the histogram.

Floderus et al. [8] gave an $O(n \log n)$-time 4-approximation algorithm to partition a 3D-histogram into a minimum number of boxes and asked whether the problem is NP-hard for histograms. We note that the NP-hardness reduction of Dielissen and Kaldewaij [3] does not hold for histograms.

Contributions: In this paper, we prove that partitioning 3D-histograms (even with height 2) into a minimum number of boxes is NP-hard (Section 3). The problem is APX-hard for 3D-histograms of height 4 (Section 4). We show that optimal partitioning must consider cuts beyond those that are “guillotine” (Section 5). We then focus on restricted classes of polyhedron (Sections 6–7). If two dimensions of the polyhedron are fixed, then we compute a minimum box partition in polynomial time. If one dimension is bounded by t, then we produce a t-approximation in polynomial time. Finally, we give a polynomial-time 2-approximation algorithm for the box partition of corner polyhedra.

Background: The problem of partitioning an orthogonal polygon into rectangles has also been considered with different objective functions, for example, minimizing the total length of the cuts (“minimum ink”) [11], avoiding very thin rectangles by minimizing the aspect ratio [16], or minimizing the so-called “stabbing number” [4]. Computing a minimum decomposition of arbitrary polygons with holes into (perhaps overlapping) convex, star-shaped, or spiral subsets is NP-hard [15]. If Steiner points are not permitted, then some partition and covering problems become polynomial-time solvable for arbitrary simple polygons [10].

2 Preliminaries

The *top surface* of a 3D-histogram H denotes the union of the faces parallel to the base, excluding the base itself. For any integer $h \geq 1$, an *h-region* of the histogram is a maximal region that all has the same face f as top surface, and f has height h.

For a point p in $\mathbb{R}^3$, we denote its x, y and z-coordinates by p_x, p_y and p_z, respectively. A set of points in H is called *independent* if there does not exist any box in H that contains two or more of these points. A set of edges in H is called a set of *forcing edges* if there does not exist any box in H that properly intersects two or more of these edges. This implies that the midpoints of the forcing edges form an independent set of points. Consequently, if there are k forcing edges in a 3D-histogram, then k is a lower bound on the size of any box partition of the histogram.

Let L be a plane parallel to one of the axis planes. We say that a *guillotine cut* along L partitions a polyhedron H into two polyhedra H_1 and H_2 if $H_1 \cup H_2 = H$ and $H_1 \cap H_2 \subset L$ (Fig. 1(c)–(d)).

Let $G = (V, E)$ be a graph with $n = |V|$ vertices and $m = |E|$ edges. We call G a *planar graph* if it admits a drawing on the Euclidean plane such that no two edges cross except possibly at a common end-point. G is *cubic* if the degree of every vertex in G is exactly three. A *vertex cover* of G is a set of vertices C in G such that for every edge (v, w), at least one of v and w belongs to C.

3 3D-Histogram Partition is NP-Hard for height ≥ 2

In this section we prove that partitioning a 3D-histogram into a minimum number of boxes (*3D-Histogram Partition*) is NP-hard even when the histogram has height two. We reduce from the problem of computing a minimum-cardinality vertex cover in a cubic planar graph, which is NP-hard [18].

Let G be a cubic planar graph and let G' be the *2-subdivision of* G, defined to be the graph obtained from G by replacing each edge of G by a path with three edges, of which the middle one is a double edge. Observe that G' is also cubic and planar. The crucial idea for constructing a histogram is to use a suitable drawing of G'. Here, an *orthogonal drawing* of a planar graph G is a planar drawing of G such that each vertex is mapped to a point in the Euclidean plane, and each edge is mapped to an axis-aligned polyline between the corresponding points. It is called *1-bend* if every polyline has exactly one bend.

Lemma 1. *Let G be a cubic planar graph. Then the 2-subdivision G' of G admits a 1-bend orthogonal drawing Γ'.*

Proof. (Sketch) Take an orthogonal drawing Γ of G with at most two bends per edge, which can be constructed in linear time [9]. For any edge without bends in Γ, apply a so-called zig-zag transformation so that it obtains exactly two bends. We obtain Γ' by replacing the drawing of each edge of e with the drawing of a path with a double edge along e's poly-line, see Fig. 2(b)–(c). □

Construction of H: Graph G' has $n' = n + 2m$ vertices and $m' = \frac{3}{2}n'$ edges. We construct a histogram H that can be partitioned into $(4n'+3m'+\alpha')$ boxes if and only if G' contains a vertex cover of size α'. Note that G' contains a vertex

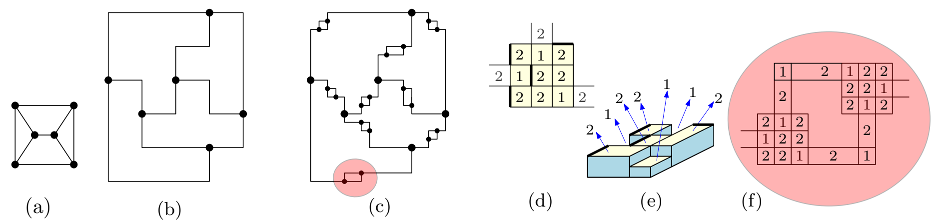

Fig. 2. (a) A cubic graph G. (b) Its orthogonal drawing Γ. (c) A 1-bend orthogonal drawing Γ' of G'. (d) Top view of a vertex gadget, where the edge connections are shown in gray. (e) Side view of a vertex gadget. (f) Connecting vertex gadgets using edge gadgets.

cover of size α' if and only if G contains a vertex cover of size $\alpha' - m$ (see e.g., see [17]), so this then proves the reduction.

We transform the 1-bend drawing Γ' of G' into the desired histogram H. Specifically, we replace each vertex v of Γ' (i.e., both original and subdivision vertices) by a vertex gadget $\lambda(v)$ (see Fig. 2(d)–(e)). The numbers in Fig. 2(d) illustrate the heights of the corresponding regions. We then replace each edge (v, w) of Γ' using an edge gadget $\lambda(v, w)$ (see Fig. 2(f)). The edge gadgets corresponding to the edges incident to v are attached to the three sides of height 1 of $\lambda(v)$. This completes the construction of H.

Let $\textsc{Opt}(H)$ be a partition of H into a minimum number of boxes. The following lemmas (whose proofs are omitted) discuss some properties of the gadgets with respect to $\textsc{Opt}(H)$. In brief, Lemma 2 follows from the forcing edges that are illustrated in Fig. 2(c). Lemmas 3–4 follow from the observation that a box that touches a face of height 2 (on the top surface) cannot touch a face of height 1, and that the edge gadgets are "non-aligned".

Lemma 2. *For every vertex gadget $\lambda(v)$, $\textsc{Opt}(H)$ contains at least four distinct boxes that lie entirely inside $\lambda(v)$. If it contains exactly four such boxes, then none of the 1-regions in $\lambda(v)$ are covered by these boxes.*

Lemma 3. *For every edge gadget $\lambda(v, w)$, $\textsc{Opt}(H)$ must contain 3 boxes that intersect $\lambda(v, w)$. No box in $\textsc{Opt}(H)$ can intersect more than one edge gadget.*

Lemma 4. *If an edge gadget $\lambda(v, w)$ is entirely covered by exactly three boxes in $\textsc{Opt}(H)$, then these three boxes cover at most one of the two 1-regions of $\lambda(v)$ and $\lambda(w)$ that are adjacent to $\lambda(v, w)$ (e.g., see Fig. 3(a)–(d)).*

Equivalence Between Instances: Given a set of r boxes, it is straightforward to verify whether the boxes are interior disjoint and cover the input histogram in polynomial time. Hence the problem 3D-Histogram Partition is in NP. Since H can be constructed in polynomial time, we can use the following lemma to obtain the NP-hardness.

Lemma 5. *G' contains a vertex cover $\mathcal{C}$ of size α' if and only if H can be partitioned into $(4n' + 3m' + \alpha')$ boxes.*

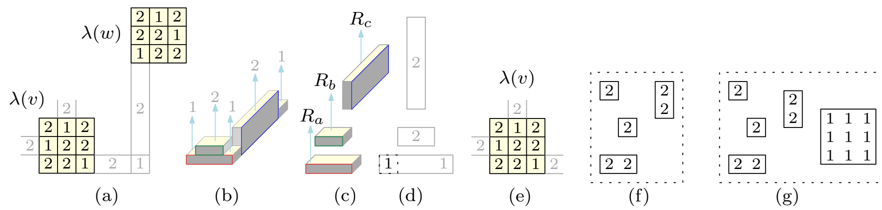

Fig. 3. (a)–(b) Illustration for $\lambda(v, w)$. (c) A partition of the edge gadget that covers the 1-region of v at (v, w). (d) A schematic representation of the partition. (e)–(g) Illustration for Lemma 5.

Proof (sketch). We construct a partition of H from a vertex cover, as follows.

A. If $v \notin \mathcal{C}$, then we use four maximal boxes to cover the 2-regions of $\lambda(v)$, as illustrated in Fig. 3(e)–(f). The remaining regions of $\lambda(v)$ are 1-regions, which will be covered by the boxes partitioning the edge gadgets.
B. If $v \in \mathcal{C}$, then we use 5 maximal boxes to cover $\lambda(v)$, e.g., see Fig. 3(g).
C. We use three maximal boxes to cover each edge gadget $\lambda(v, w)$. Note that either v or w must lie in $\mathcal{C}$. If $v \notin \mathcal{C}$, then one of these boxes will cover the 1-region of $\lambda(v)$ at $\lambda(v, w)$, e.g., see Fig. 3(a)–(d). Similarly, if $w \notin \mathcal{C}$, then one of these boxes will cover the 1-region of $\lambda(w)$ at $\lambda(v, w)$.

One easily verifies that Steps A–C partition the histogram H into $(3m' + 4n' + \alpha')$ boxes. For the other direction, assume that H admits a partition $\mathcal{B}$ with $(3m' + 4n' + \alpha')$ boxes, and construct a vertex cover of size at most α' in G'. By Lemma 2, every vertex gadget contributes to at least four distinct boxes in $\mathcal{B}$, which corresponds to the 2-regions. Hence we use at least $4n'$ boxes of $\mathcal{B}$ to cover those regions. Note that all these boxes lie entirely inside the vertex gadgets. By Lemma 3, every edge gadget must use at least three distinct boxes, which altogether sum up to at least $3m'$. These boxes may also cover some 1-regions of the vertex gadgets (e.g., see Lemma 4). Since $\mathcal{B}$ contains at most $(3m' + 4n' + \alpha')$ boxes, we have at most α' boxes remaining to cover the remaining 1-regions of the vertex gadgets. We now construct a set $\mathcal{S}$ as follows: (a) If a vertex gadget $\lambda(v)$ contains more than four boxes lying entirely inside $\lambda(v)$, then we include v into $\mathcal{S}$. (b) If four or more boxes intersect an edge gadget $\lambda(v, w)$, then we choose one of v and w arbitrarily into $\mathcal{S}$.

Note that each step can be charged uniquely to one of the remaining α' boxes, i.e., the box charged in Step (a) lies entirely inside $\lambda(v)$ and hence cannot be charged again in Step (b). Hence the number of vertices in $\mathcal{S}$ is at most α'. One argues that $\mathcal{S}$ is a vertex cover based on the observation that either (a) or (b) must apply for each edge gadget. ☐

Theorem 1. *Partitioning a 3D-histogram into a minimum number of boxes is NP-hard, even when the histogram is of height two.*

4 3D-Histogram Partition is APX-Hard for height ≥ 4

In this section we prove that partitioning a 3D-histogram into a minimum number of boxes is APX-hard, even for histograms of height 4. We reduce from the problem of computing a minimum-cardinality vertex cover in a cubic graph (not necessarily planar), which is APX-hard [1].

Construction of H: Let $V = \{v_1, \ldots, v_n\}$ and $E = \{e_1, \ldots, e_m\}$ be the vertices and edges of G. Consider an integer grid of size $(8n+1) \times (2n+10m+1)$. Column $8i-3$ is assigned to v_i, and column $8i+1$ is assigned to the transition from v_i to v_{i+1} (we write $w_{i,i+1}$ for short as in Fig. 4(b)). Below the grid we add a staircase that descends at each column of v_i or $w_{i,i+1}$, and here add a *tooth*, i.e., a square for which all but the top sides are on the boundary. Edge e_j is assigned to row $10i+2n+4$. For each edge, cut out H-shaped holes in the polygon where the row of the edge meets the columns of its endpoints. These holes have width 7, height 7 or more, and remove four (five) squares from the column of the left (right) endpoint. The resulting polygon P (Fig. 4(a)) forms the base of the histogram (except for some additions via edge gadgets that will be listed below). Extrude all of P to height 2; we call the result the *platform*.

For each edge $e = (v, w)$, we add an *edge gadget* $\lambda(v, w)$ that consists of two *endpoint gadgets* and a *connector*. Here, the endpoint gadget at v consists of the four (five) squares from the vertex column of v (we call these the *decision column* $\lambda(v, e)$) as well as a surrounding polygon; all of these have height 1 (see Fig. 4(d)–(e)). The *connector* connects the two endpoint gadgets via a sequence of 2-regions, 3-regions and 4-regions along the row of the edge; it sits partially on the endpoint gadgets and partially on the platform (see Fig. 4(f)–(g)). This completes the construction of H.

To argue the correctness, we fix a set F of $13n + 1$ edges (shown in bold in Fig. 4(b)) that can easily be seen to be forcing edges. We use the $2n + 1$ horizontal top edges of the teeth, as well as the $2n$ horizontal top edges that lie between these teeth. For each edge, we select 6 further horizontal edges from the hole boundaries of its two endpoints, see Fig. 4(b); we assume that these hole boundaries were chosen that none of them have the same y-coordinate. Consequently, we obtain a set of $(2n + 1) + (2n) + (6m) = 4n + 9n + 1$ forcing edges, all of which are on top of the platform.

We must argue how many boxes are needed to cover most (but not all) of an edge gadget $\lambda(v, w)$. We say that a *sub-partition of* $\lambda(v, w)$ is a set of disjoint boxes that cover the entire edge gadget except that they may leave one or both decision columns $\lambda(v, e)$ or $\lambda(w, e)$ uncovered. Given a partition of the histogram, we *charge* a box B to edge-gadget $\lambda(v, w)$ if either B intersects the interior of $\lambda(v, w)$ or if B lies inside the platform and the top front edge of B lies on the boundary of the 4-region of $\lambda(v, w)$. (In particular, such a box cannot cover an edge in F and it can only be charged to one edge-gadget.) Crucial for the reduction is the following lemma:

Lemma 6. *The edge gadgets satisfy the following properties: (P_1) Every sub-partition of an edge gadget has least 12 boxes charged to it. (P_2) Every partition*

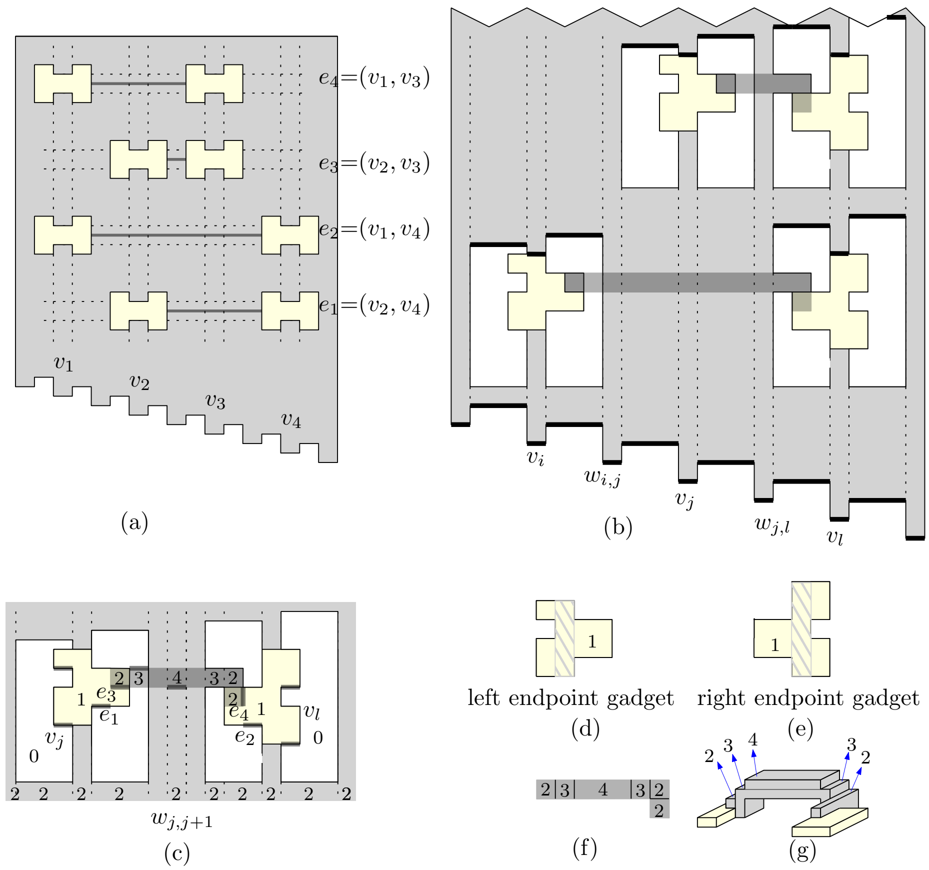

Fig. 4. (a) Schematic representation and (b) top view of H. (c) Illustration for the heights near one edge gadget. (d)–(e) Endpoint gadgets and decision columns. (f)–(g) A connector gadget is shown in gray.

of an edge gadget (covering both decision columns) has at least 13 boxes charged to it.

Proof (sketch). Consider a (sub-)partition $\mathcal{B}$ of some edge gadget $\lambda(v,w)$. Fig. 4(c) shows some forcing edges of $\lambda(v,w)$. Of these, there are three each in the endpoint gadgets, forcing three boxes each, and four more in the connector-gadget.

In fact, the connector-gadget requires five boxes to be charged, as can be seen as follows: If no box of $\mathcal{B}$ intersects the platform below the connector-gadget, then we can find another independent point (on the downward-facing face of the 3-region on the right in Fig. 4(c), at height 1.5). If some box B of $\mathcal{B}$ intersects the platform, then some other box of the partition of $\mathcal{P}$ must share a side with B, and this other box is also charged to $\lambda(v,w)$.

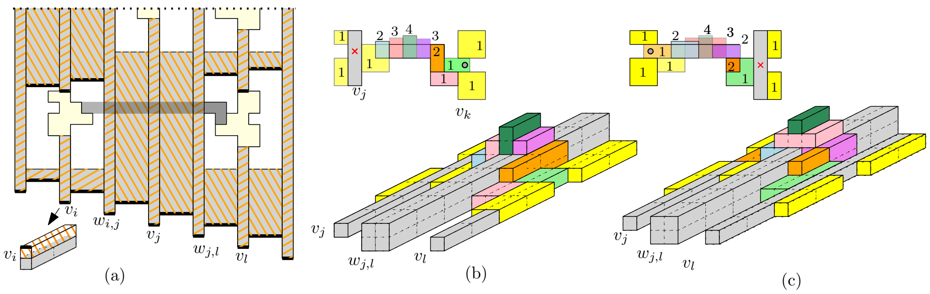

Fig. 5. (a) The maximal boxes determined by the forcing edges. (b)–(c) Partition of the endpoint gadget excluding the left and right decision column, respectively.

Thus we have now 11 boxes to be charged to $\lambda(v, w)$. One can also easily verify that if at one connector-gadget the decision column is covered, then this requires one additional box not counted elsewhere. This proves the claim in all cases except the one where neither decision-column is covered, each endpoint-gadget uses exactly three boxes, and the connector-gadget has exactly five boxes charged to it. One can verify that this is impossible if none of the boxes overlap. □

Equivalence Between Instances: We now prove that H can be partitioned into $(|F| + 12m + k)$ boxes if and only if G has a vertex cover of size at most k.

Lemma 7. *If G has a vertex cover $\mathcal{C}$ of size k, then H admits a partition into $(|F| + 12m + k)$ boxes.*

Proof (sketch). For each vertex $v \in \mathcal{C}$, construct a box that has width and height 1 and whose depth is so large that is spans the entire column of v. In particular it covers all decision columns $\lambda(v, e)$ of incident edges of v, e.g., see the gray box for v_i in Fig. 5(a). The rest of the platform can be covered using one box per forcing edge. Finally, there are sub-partitions of an edge gadget $\lambda(v, w)$ using 12 boxes (e.g., Fig. 5(b)–(c)) so that $\lambda(v, e)$, $\lambda(w, e)$, or neither, is covered. Applying the suitable one to each edge (depending on whether $v \in \mathcal{C}$, $w \in \mathcal{C}$, or $v, w \in \mathcal{C}$) gives the desired partition. □

Lemma 8. *If H admits a partition $\mathcal{B}$ into $(|F| + 12m + k)$ boxes, then G has a vertex cover of size at most k.*

Proof. We construct a vertex cover $\mathcal{C}$ of G as follows: (a) For every decision column $\lambda(v, e)$, if the box covering it lies entirely in the column of v, then add v to $\mathcal{C}$. (b) For an edge (v, w), if neither v nor w belongs to $\mathcal{C}$, and at least 13 boxes are charged to $\lambda(v, w)$, then arbitrarily add one of v and w to $\mathcal{C}$.

We first show that $\mathcal{C}$ contains at most k vertices. The set F of forcing edges determines $|F|$ boxes that we denote by R_F. All of them cover no decision column and cover no part of an edge-gadget or are charged to it. $\mathcal{B} - R_F$ has $12m + k$

boxes. By Lemma 6 at least 12 of them are charged to each edge. This leaves at most k boxes that lead to an addition to $\mathcal{C}$.

Suppose now for a contradiction that for some edge $e = (v, w)$ neither v nor w belongs to $\mathcal{C}$. Then the boxes covering $\lambda(v, e)$ and $\lambda(w, e)$ cannot lie entirely inside their corresponding vertex columns. In other words, Step (a) did not apply. In this scenario, both these decision columns are covered using boxes from the edge-gadget, so by Lemma 6 at least 13 boxes are charged to $\lambda(v, w)$. Hence by Step (b), either v or w must belong to $\mathcal{C}$. $\square$

Theorem 2. 3D-Histogram-Partition *is APX-hard.*

Proof. Let $\mathcal{C}^*$ and S^* be the optimum vertex cover of G and the optimum box partition of H, respectively. Assume that we had an $(1 + \epsilon)$-approximation algorithm for 3D-Histogram-Partition that computes a solution S from which we extract a vertex cover $\mathcal{C}$. By Lemmas 7–8, we have $|S| = |F| + 12m + |\mathcal{C}|$ and $|S^*| = |F| + 12m + |\mathcal{C}^*|$. Since G is cubic, $|\mathcal{C}^*| \geq n/3$. Finally observe that $|S^*| \leq |F| + 12m + |\mathcal{C}^*| \leq 13n + 1 + 18n + (n - 1) = 32n$. Hence we get

$$\frac{|\mathcal{C}|}{|\mathcal{C}^*|} = \frac{|S|-|F|-12m}{|S^*|-|F|-12m} \leq \frac{(1+\epsilon)|S^*|-|F|-12m}{|S^*|-|F|-12m} = 1+\frac{\epsilon|S^*|}{|\mathcal{C}^*|} \leq 1+\frac{32n\epsilon}{n/3} = 1+96\epsilon,$$

implying an approximation algorithm for vertex cover. The APX-hardness of 3D-Histogram-Partition now follows from the APX-hardness of minimum vertex cover. $\square$

5 Partitioning Using Guillotine Cuts

In this section we show that there is an infinite family of 3D-histograms that cannot be optimally partitioned using guillotine cuts, whereas 2D-polygons can be partitioned optimally using such cuts by first cutting along “good diagonals” [5].

We say that $\mathcal{P}$ is a partitioning of a polyhedron H into boxes using guillotine cuts if $\mathcal{P}$ is a partition of H into boxes and there is a sequence $\mathcal{P}_0 = \{H\}, \mathcal{P}_1, \ldots, \mathcal{P}_k = \mathcal{P}$ of sets of polyhedra such that every $\mathcal{P}_{i+1}$ is obtained from $\mathcal{P}_i$ by partitioning $\mathcal{P}_i$ using guillotine cuts.

Theorem 3. *There is an infinite family of 3D-histograms that cannot be partitioned optimally using only guillotine cuts.*

Proof. We first refer the reader to the histogram H of Fig. 6(a). Since H has five faces of distinct height, any partition of H into boxes would require 5 boxes, and H admits such a partition (e.g. by cutting along the edges of the top view in Fig. 6(b)). We now show that H cannot be partitioned into 5 boxes if we restrict the cuts to be guillotine.

Consider starting with a vertical guillotine cut, i.e., a cut perpendicular to the x-axis or y-axis. Any such cut results in two polyhedra: one with at least 4 faces of distinct height, and another with at least 2 faces of distinct height. Any further cutting of these polyhedra will result in at least 6 boxes, a contradiction.

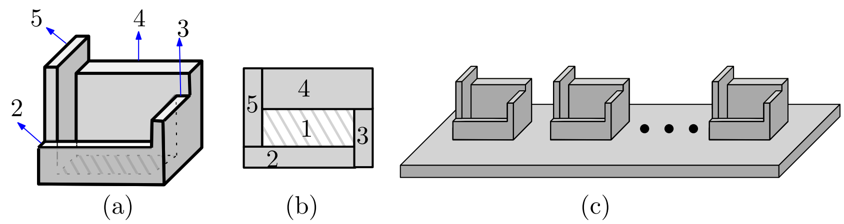

Fig. 6. (a) Illustration for H. (b) The top view of H. (c) Construction of H_k.

If instead we start with a horizontal cut (perpendicular to the z-axis), we have 4 choices: cutting at heights 1, 2, 3, or 4. Cutting at any height other than 1 results in two polyhedra, one of which with 5 boxes in its optimal partitioning, a contradiction. Assume that we start by cutting at height 1. Any subsequent vertical guillotine cut results in two polyhedra: one with at least 3 faces of distinct height, and another with at least 2. Therefore, no vertical guillotine cuts are acceptable. Any subsequent horizontal cut immediately result in a contradiction: one of the polyhedra resulting from such a horizontal guillotine cut has 4 boxes in its optimal partitioning, and together with the 1 box from the first cut and the (at least) 1 other box from this cut we have at least 6 boxes.

It is now straightforward to create an infinite family of polyhedra $H_1(= H), H_2, \ldots, H_k$ by attaching up to k copies of H on a rectangular box, as in Fig. 6(c), such that none of them can be partitioned optimally using only guillotine cuts. $\square$

6 Orthogonal Polyhedra with Bounded Dimensions

In this section we focus on orthogonal polyhedra with bounded dimensions (recall that all vertex-coordinates are assumed to be integers). If one dimension of the input polyhedron $\mathcal{P}$ is bounded by t, then we construct a t-approximate box partition. If two dimensions are bounded, then we compute an optimal box partition in polynomial time. This works for all polyhedra of bounded dimensions, even if they are not histograms or have holes.

So let $\mathcal{P}$ be a polyhedron that resides within the $[0, W] \times [0, L] \times [0, H]$ box. If H is bounded by t, then we partition the polyhedron into (up to) t sets of polyhedra $\mathcal{P}_i$ where $0 \le i < t$ and $\mathcal{P}_i$ is bounded by the planes $z = i, z = i + 1$. For each $\mathcal{P}_i$, we compute an optimal box partition $B_i^* = \textsc{Opt}(\mathcal{P}_i)$ using the algorithm for partitioning 2D-polygons [5]. We claim that $\bigcup_i B_i^*$ gives a partition that is within a factor of t of the optimum.

Fix an optimal partition $\textsc{Opt}(\mathcal{P})$ and partition it by the planes $z = i$. Let B_i be the boxes between the planes $z = i$ and $z = i + 1$. Then $|B_i| \ge |B_i^*|$ but also $|B_i| \le |\textsc{Opt}(\mathcal{P})|$. Hence $\sum_i |B_i^*| \le \sum_i |B_i| \le t \cdot |\textsc{Opt}(\mathcal{P})|$ and we have:

Lemma 9. *For orthogonal polyhedra with one dimension bounded by t, a minimum box partition can be approximated within a factor of t in polynomial time.*

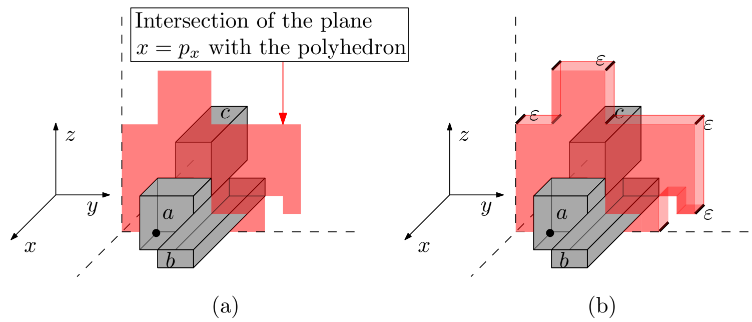

Fig. 7. Illustration for Lemma 10.

Consider now the scenario when two dimensions are fixed, e.g., $W \cdot L \in O(1)$. We rely on the following lemma.

Lemma 10. *Any orthogonal polyhedron $\mathcal{P}$ with vertices having integer coordinates can be optimally partitioned into boxes where the coordinates are integral.*

Proof (sketch). Assume that in a partition some box-face is within a plane (say plane $x = p_x$) for which p_x is not integral. Then we can shift all box-boundaries within that plane to lie within $x = p_x + \varepsilon$ instead (see Fig. 7) to get closer to a solution where all coordinates are integral. $\square$

Let the *voxel* $v_{i,j,k}$ be the unit cube $[i-1, i] \times [j-1, j] \times [k-1, k]$ and let the *column* $c_{i,j}$ be all the voxels $\{v_{i,j,k} : 1 \le k \le H\}$. (In the following, whenever the range of i, j is unspecified then we mean $1 \le i \le W$ and $1 \le j \le L$.) We have $W \cdot L \in O(1)$ columns, and hence $O(H)$ voxels. Consider some box partition $\mathcal{B}$ of $\mathcal{P}$ that uses only boxes with integer coordinates. Let $\mathcal{B}'$ be some set of boxes obtained from $\mathcal{B}$ by removing (repeatedly) some box whose entire top is visible to infinity, or becomes visible after some other boxes in $\mathcal{B} - \mathcal{B}'$ were removed. Boxes $\mathcal{B}'$ cover a subset $\mathcal{P}'$ of $\mathcal{P}$, and this subset can be described as follows: For each column $c_{i,j}$, $\mathcal{P}'$ contains all voxels of $c_{i,j}$ that belong to $\mathcal{P}$, up to some limit $b_{i,j}$, and then contains no other voxels of $c_{i,j}$.

This observation is the key idea for a dynamic programming algorithm to find the optimum partition of $\mathcal{P}$. For any integer values $b_{i,j}$, define the polyhedron $\mathcal{P}[\{b_{i,j}\}_{i,j}]$ to be the polyhedron obtained from $\mathcal{P}$ by removing for each column $c_{i,j}$ all voxels $v_{i,j,q}$ with $q > b_{i,j}$. For each such polyhedron, we compute (recursively) the size $T[\{b_{i,j}\}_{i,j}]$ of the optimal box partition. This gives the optimal box partition for $\mathcal{P} = \mathcal{P}[\{H\}_{i,j}]$.

We can fill an entry $T[\{b_{i,j}\}_{i,j}]$by considering any box $B = [i_1, i_2] \times [j_1, j_2] \times [k_1, k_2]$ that could be part of a box partition of $\mathcal{P}[\{b_{i,j}\}_{i,j}]$ such that the top face of B is visible to infinity. In particular, we must have $b_{i,j} = k_2$ and $v_{i,j,k} \subset \mathcal{P}$ for all $i_1 < i \le i_2, j_1 < j \le j_2$, and $k_1 < k \le k_2$. If this is satisfied, then one possible value for $T[\{b_{i,j}\}_{i,j}]$ is to add one to the value for $T[\{b'_{i,j}\}_{i,j}]$ (where $b'_{i,j} = k_1$ for all $i_1 < i \le i_2, j_1 < j \le j_2$ and $b'_{i,j} = b_{i,j}$ otherwise).

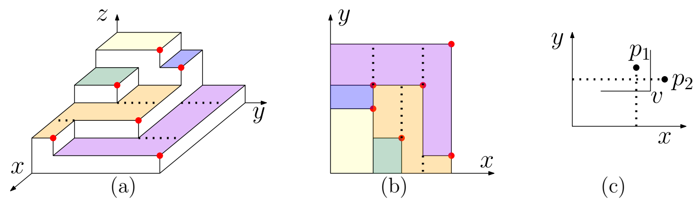

Fig. 8. (a) $\mathcal{P}$, with peaks shown as red dots. (b) $\mathcal{H}$, the projection of $\mathcal{P}$ to the $z = 0$ plane, and a partition of $\mathcal{H}$ into rectangles. (c) A forbidden corner in $\mathcal{H}$.

There are $W^2L^2 \in O(1)$ possibilities for i_1, i_2, j_1, j_2, and for each of them, we can find the only possible value k_2 in $O(W \cdot L) = O(1)$ time and all possible values of k_1 in $O(H)$ time. One update to the table can hence be done in $O(H)$ time. There are $O(H^{WL})$ table-entries, so the entire dynamic program takes $O(H^{O(1)})$ time. We may assume that any plane $z = i$ for integral i contains at least one vertex of $\mathcal{P}$ (else we could shrink the polyhedron to obtain a combinatorially equivalent one) so that $H \in \Theta(n)$. Therefore we can find the optimal partition in $O(n^{O(1)})$ time.

Theorem 4. *Given an orthogonal polyhedron $\mathcal{P}$ such that two dimensions of $\mathcal{P}$ are bounded, one can compute a minimum box partition of $\mathcal{P}$ in polynomial time.*

7 Corner Polyhedra

In this section we give a polynomial-time algorithm with approximation factor 2 for partitioning a *corner polyhedron* into a minimum number of rectangular boxes. This improves on the approximation factor of 4 for histograms.

A *corner polyhedron* (as defined by Eppstein [6]) is an orthogonal polyhedron in which all but three "back" faces are oriented towards the vector $(1, 1, 1)$, i.e., visible from a point at infinity on the line $x = y = z$. See Fig. 8. Without loss of generality we will assume that the three back faces intersect at the vertex $(0, 0, 0)$. A corner polyhedron can be drawn in the plane by isometric projection with all vertices except $(0, 0, 0)$ visible. For any point $p = (p_x, p_y, p_z)$ inside a corner polyhedron, the three orthogonal line segments connecting p to the planes $z = 0$, $y = 0$, and $x = 0$ are contained in the polyhedron. This implies that a corner polyhedron is a histogram with any of the three back faces as the base.

A *peak* of a corner polyhedron is a vertex that is a local maximum in the direction $(1, 1, 1)$. Equivalently a peak is a vertex where the solid angle is $4\pi/8$, not including the vertices that lie on the axis planes. Let k be the number of peaks. Observe that the peaks form a set of independent points, thus k is a lower bound on the number of boxes needed in a partition. We will show that any corner polyhedron can be partitioned into $2k$ boxes, which gives the approximation factor 2.

Lemma 11. *A corner polyhedron $\mathcal{P}$ with k peaks can be partitioned into $2k$ boxes. Furthermore, such a partition can be found in polynomial time.*

Proof. Project $\mathcal{P}$ onto one of the axis planes, say $z = 0$. Call the result $\mathcal{H}$. Then $\mathcal{H}$ is a histogram partitioned into orthogonal polygons that correspond to the top faces of $\mathcal{P}$.

We claim that each such polygon consists of two *monotone* chains, each consisting of edges in the $+x$ and $-y$ directions. See Fig. 8(b). To justify this claim, observe that if a polygon of $\mathcal{H}$ were not monotone then it would have a vertex v where one edge goes to the left and one edge goes up, as shown in Fig. 8(c). Let p_1 be a point just above and left of v, and let p_2 be a point just above and right of v. Assume that p_1 and p_2 are the projections of points q_1 and q_2 on the surface of $\mathcal{P}$. We consider their z coordinates. If $z(q_1) > z(q_2)$ then a line from q_1 to the $y = 0$ plane leaves $\mathcal{P}$, and if $z(q_1) < z(q_2)$ then a line from q_2 to the $x = 0$ plane leaves $\mathcal{P}$—a contradiction in both cases.

Given that the polygons of $\mathcal{H}$ are monotone, we can partition them into rectangles by adding, from each peak vertex in $\mathcal{H}$, two vertical (i.e., in the y direction) segments, one going upwards and one going downwards. Each line segment stops when it is blocked by the interior of a horizontal edge of $\mathcal{H}$. This partitions $\mathcal{H}$ into at most $2k$ rectangles. Expand each rectangle into a 3D box from $z = 0$ to the maximum possible height. The result is $2k$ boxes that partition $\mathcal{P}$. See the final partition of $\mathcal{P}$ in Fig. 8(a). □

Theorem 5. *There is a polynomial-time 2-approximation algorithm to partition a corner polyhedron into boxes.*

8 Open Problems

1. Is there a constant-factor approximation algorithm for the case of general 3D orthogonal polyhedra?
2. For histograms, there is a 4-approximation algorithm [8]. Can the approximation factor of 4 be reduced?
3. For corner polyhedra, we gave a 2-approximation algorithm. Is there a PTAS, or even a polynomial-time algorithm?
4. What about other special cases of histograms, for example "convex polyhedra"[3] and "orthoballs"[2]?

For all these questions, the concept of independent points may be useful. While there are histograms for which the optimal box partition has a higher cardinality than any independent set of points (e.g. the one in Fig. 1(c) has only three independent points but requires four boxes), the maximum cardinality of an independent set of points could serve as a lower bound for an approximation algorithm. Can it be computed efficiently?

Acknowledgements: This work was done as part of the Algorithms Problem Session at the University of Waterloo. We thank the other participants for valuable discussions. Research of T.B. and A.L. supported by NSERC, M.D. supported by Vanier CGS, M.D. and D.M. supported by an NSERC PDF.

References

1. P. Alimonti and V. Kann. Some APX-completeness results for cubic graphs. *Theoretical Computer Science*, 237(1-2):123–134, 2000.
2. F. Barrera-Cruz, T. C. Biedl, M. Derka, S. Kiazyk, A. Lubiw, and H. Vosoughpour. Turning orthogonally convex polyhedra into orthoballs. In *Proc. of CCCG*, 2014.
3. V. J. Dielissen and A. Kaldewaij. Rectangular partition is polynomial in two dimensions but NP-complete in three. *Inf. Process. Lett.*, 38(1):1–6, 1991.
4. S. Durocher and S. Mehrabi. Computing conforming partitions of orthogonal polygons with minimum stabbing number. *Theor. Comput. Sci.*, 689:157–168, 2017.
5. D. Eppstein. Graph-theoretic solutions to computational geometry problems. In *Proc. of WG*, volume 5911 of *LNCS*, pages 1–16, 2009.
6. D. Eppstein and E. Mumford. Steinitz theorems for simple orthogonal polyhedra. *Journal of Computational Geometry*, 5(1):179–244, 2014.
7. L. Ferrari, P. V. Sankar, and J. Sklansky. Minimal rectangular partitions of digitized blobs. *Computer Vision, Graphics, and Image Processing*, 28(1):58–71, 1984.
8. P. Floderus, J. Jansson, C. Levcopoulos, A. Lingas, and D. Sledneu. 3D rectangulations and geometric matrix multiplication. *Algorithmica*, 2016. In press.
9. G. Kant. Drawing planar graphs using the canonical ordering. *Algorithmica*, 16(1):4–32, 1996.
10. M. Keil and J. Snoeyink. On the time bound for convex decomposition of simple polygons. *International Journal of Computational Geometry & Applications*, 12(03):181–192, 2002.
11. A. Lingas, R. Pinter, R. Rivest, and A. Shamir. Minimum edge length partitioning of rectilinear polygons. In *Proc. of Annual Allerton Conf. on Communication, Control, and Computing*, volume 10, pages 53–63, 1982.
12. W. Lipski. Finding a Manhattan path and related problems. *Networks*, 13(3):399–409, 1983.
13. W. Lipski, E. Lodi, F. Luccio, C. Mugnai, and L. Pagli. On two-dimensional data organization II. *Fundamenta Informaticae*, 2:245–260, 1979.
14. T. Ohtsuki. Minimum dissection of rectilinear regions. In *Proc. IEEE Int. Symp. Circuits and Systems*, pages 1210–1213, 1982.
15. J. O'Rourke and K. J. Supowit. Some NP-hard polygon decomposition problems. *IEEE Trans. Information Theory*, 29(2):181–189, 1983.
16. J. O'Rourke and G. Tewari. The structure of optimal partitions of orthogonal polygons into fat rectangles. *Computational Geometry*, 28(1):49–71, 2004.
17. S. Poljak. A note on stable sets and colorings of graphs. *Commentationes Mathematicae Universitatis Carolinae*, 15(2):307–309, 1974.
18. R. Uehara. NP-complete problems on a 3-connected cubic planar graph and their applications. *Tech. Report TWCU-M-0004, Tokyo Woman's Christian Univ.*, 1996.