Mapping Multiple Regions to the Grid with Bounded Hausdorff Distance

Ivor van der Hoog, Mees van de Kerkhof, Marc van Kreveld, Maarten Löffler, Frank Staals, Jérôme Urhausen(✉), and Jordi L. Vermeulen

Department of Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands
j.e.urhausen@uu.nl

Abstract. We study a problem motivated by digital geometry: given a set of disjoint geometric regions, assign each region R_i a set of grid cells P_i, so that P_i is connected, similar to R_i, and does not touch any grid cell assigned to another region. Similarity is measured using the Hausdorff distance. We analyze the achievable Hausdorff distance in terms of the number of input regions, and prove asymptotically tight bounds for several classes of input regions.

Keywords: Computational geometry · Digital geometry · Hausdorff distance · Simple polygons

1 Introduction

Digital geometry is concerned with the proper representation of geometric objects and their relationships using a grid of pixels. This greatly simplifies both representation and many operations, but the downside is that common properties of geometric objects no longer hold. For example, it may be that two digitized lines intersect in multiple connected components. One objective of digital geometry is how to *consistently* digitize a set of geometric objects. Another objective is the presentation of vector objects with bounded error, using subsets of pixels.

Early results in digital geometry were mostly concerned with consistency and arose in computer vision. For a survey, see Klette and Rosenfeld [11,12]. More recently, also error bounds under the Hausdorff distance have been studied. Chun et al. [5] investigate the problem of digitizing rays originating in the origin to digital rays such that certain properties are satisfied. They show that rays can be represented on the $n \times n$ grid in a consistent manner with Hausdorff distance $O(\log n)$. This bound is tight in the worst case. By ignoring one of the consistency conditions, the distance bound improves to $O(1)$. Their research is extended by Christ et al. [3] to line segments (not necessarily starting in the origin), who obtain the logarithmic distance bound in this case as well. A possible extension to curved rays was developed by Chun et al. [4]. Other results with a digital geometry flavor within the algorithms community are those on snap rounding [6,7,10], integer hulls [1,9], and discrete schematization [13].

A. Lubiw et al. (Eds.): WADS 2021, LNCS 12808, pp. 627–640, 2021.
https://doi.org/10.1007/978-3-030-83508-8_45

In a recent paper, Bouts et al. [2] showed that any simple polygon, no matter how detailed, can be represented by a simply connected set of unit pixels such that the Hausdorff distance to and from the input is bounded by $3\sqrt{2}/2$.

Contribution. We extend the result from [2] to multiple regions, see Fig. 1. We investigate several restrictions on the class of regions and we show that stricter restrictions allow for pixel representations with a smaller symmetric Hausdorff distance. All our bounds are tight. We express our bounds in the number of input regions. Our results are shown in Table 1; they are fundamental results on the error that may be incurred when converting vector to grid representations, a common operation in computer graphics and GIS.

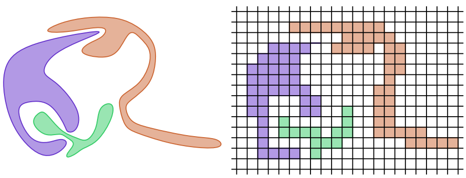

Fig. 1. Three disjoint simply connected regions and a representation by simply connected sets of disjoint pixels.

We do not make any assumptions on the resolution of the input. If the minimum distance between any pair of polygons is at least some constant (e.g., $4\sqrt{2}$ is enough), then we can realize a constant Hausdorff bound in all cases by applying the results from Bouts et al. [2] separately on each polygon. We consider the case where no such assumptions are made.

Table 1. Worst-case bounds on Hausdorff distances for m regions; β is constant.

Region class	Points	Convex β-fat	Convex	Two regions	Three regions
Hausdorff distance	$\Theta(\sqrt{m})$	$\Theta(\sqrt{m})$	$\Theta(m)$	$\Theta(1)$	unbounded

Notation and Definitions. We denote by Γ the (infinite) unit grid, whose unit squares are referred to as *pixels*. The *(symmetric) Hausdorff distance* between two sets $A, B \subset \mathbb{R}^2$ is defined as $H(A,B) = \max\{\max_{a\in A}(\min_{b\in B}(|ab|)), \max_{b\in B}(\min_{a\in A}(|ab|))\}$, where $|ab|$ is the distance between the points a and b. Further we denote by $H'(A,B) = \max\{H(A,B), H(\partial A, \partial B)\}$ the (symmetric) Hausdorff distance between the sets themselves and between their boundaries. See Fig. 2 for an example where the distinction between $H(\cdot,\cdot)$ and $H'(\cdot,\cdot)$ is important.

Let $\mathcal{R} = \{R_1, R_2, \ldots R_m\}$ be a set of m disjoint simply connected regions in the plane. In this paper, we show how to assign a subset of the pixels Two such *grid polygons* are disjoint if they do not meet in any edge or vertex of the grid. A grid polygon is connected $P_i \subset \Gamma$ to each region $R_i \in \mathcal{R}$, such that the result is a set of m disjoint simply connected regions. if its

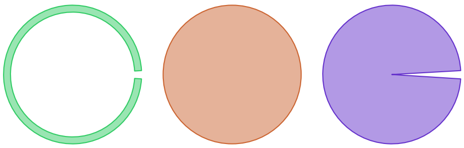

Fig. 2. The Hausdorff distance between the green and red regions is large while the Hausdorff distance between their boundaries is small. The inverse is true for the red and purple regions. (Color figure online)

pixels are connected by edge adjacency, and simply connected if it is connected and its complement is also connected by edge adjacency. We call the set $\{P_1, P_2, \ldots, P_m\}$ of such grid polygons a *valid assignment* for $\mathcal{R}$.

Overview. We are interested in finding for any set of regions $\mathcal{R}$ a valid assignment such that for all i the (symmetric) Hausdorff distance between R_i and P_i is at most h, and the (symmetric) Hausdorff distance between their boundaries is also at most h. In general, a worst-case bound on h will be a function of m. We study this problem under several restrictions on $\mathcal{R}$; refer to Table 1. For each class of restrictions, first we show that there is a set of regions in that class for which any valid assignment contains at least one region R_i with a grid polygon P_i where $H'(R_i, P_i) = \Omega(h)$. Second we show that for any set of regions in that class, we can find a valid assignment such that for all regions $R_i \in \mathcal{R}$ with corresponding grid polygon P_i, we have $H'(R_i, P_i) = O(h)$. Hence, our bounds are asymptotically tight.

We may interpret a solution to our problem as a *coloring* of Γ: each pixel $q \in \Gamma$ is assigned one color in $C = \{c_1, \ldots c_m\} \cup \{b\}$, where c_i is the color of the input region R_i and b is the background color.

Our upper bound constructions all follow a similar scheme. Let Γ_k be a coarsening of the grid Γ whose cells have $k \times k$ pixels. We call these cells *superpixels.* We will determine for each region from $\mathcal{R}$ which superpixels it contains and which ones it properly intersects. If a region R_i contains a superpixel, then all pixels of Γ in that superpixel will be part of P_i. If R_i properly intersects a superpixel, we ensure that at least one, but not all pixels in that superpixel will be part of P_i. A superpixel not intersecting R_i will have no pixels in P_i. The main challenge is then finding a scheme by which each grid polygon becomes simply connected yet all remain disjoint. It is then relatively straightforward to see that $H'(R_i, P_i) \leq k\sqrt{2}$.

2 Input Regions are Points

In this section we first consider the simplest possible case, namely, $\mathcal{R}$ is a set of points. We will construct a map that assigns points to pixels such that the symmetric Hausdorff distance between each point and its corresponding pixel is bounded. For a lower bound, consider a set of m points $\mathcal{R}$ that all lie within a single pixel. If we want to assign each point to a unique pixel, we clearly need to use m different pixels. Any set of m pixels has diameter $\Omega(\sqrt{m})$, so at least one of the point regions will be mapped to a pixel at distance $\Omega(\sqrt{m})$.

We now present a scheme that maps any set of m points $\mathcal{R}$ to a set of pixels, such that the symmetric Hausdorff distance between any point and its pixel is at most $O(\sqrt{m})$. Let Γ_k be a coarsening of Γ with $k = 2\lceil\sqrt{m}\rceil$. Associate each region in $\mathcal{R}$ with the superpixel that contains it. Each superpixel has the space to accommodate m disjoint pixels without using the bottom row and right column by using exactly the odd numbered rows and columns. Any assignment of the points to these pixels is easily seen to have Hausdorff distance $O(\sqrt{m})$.

Theorem 1. *If $\mathcal{R}$ is a set of m points, a valid assignment exists such that for each region $R_i \in \mathcal{R}$ with a corresponding region P_i, we have $H'(R_i, P_i) = O(\sqrt{m})$. Furthermore, there exists a set $\mathcal{R}$ of m points such that for every valid assignment we have $H(R_i, P_i) = \Omega(\sqrt{m})$.*

3 Input Regions are Convex β-fat Regions

A connected region R is β-fat if for some point t in R, the ratio of the radius of the smallest t-centered circle containing R and the radius of the largest t-centered circle contained in R, is β (or larger) [14]. Observe that the only regions that are 1-fat are points and disks, as points are β-fat regions for *any* $\beta \geq 1$ by convention. In this section we consider the class $\mathcal{R}$ of convex β-fat regions for a constant β. From Sect. 2 it follows that for any m, there exists a set of m regions for which the (symmetric) Hausdorff distance between $\mathcal{R}$ and any valid assignment is $\Omega(\sqrt{m})$.

Let $\mathcal{R}$ be a set of convex β-fat regions and let Γ_k be a coarsening of Γ with $k = 2\lceil\sqrt{m}\rceil + 3$. We present an algorithm that maps $\mathcal{R}$ to a set of grid polygons $\mathcal{P}$, such that the symmetric Hausdorff distance between any region R_i and its assigned region P_i is at most $O(\beta\sqrt{m})$.

Lemma 1. *Let R be a convex β-fat region, and let p be a point in R. Either R has diameter less than $16\beta k$, or R contains a superpixel within distance $16\beta k$ from p.*

This leads to the following algorithm with two cases for each region R_i, depending on the set of superpixels $\mathcal{S}_i$ contained in R_i.

Case 1: $\mathcal{S}_i$ is empty. We select any superpixel S intersected by R_i and we assign R_i to a unique pixel in S while using neither the topmost, bottommost, leftmost, or rightmost rows and columns, similar to the procedure in Sect. 2. This pixel has a distance of at most $16\beta k + \sqrt{2}k$ to any point on R_i since R_i has diameter smaller than $16\beta k$ by Lemma 1. This also means that for each such region R_i, we have $H'(R_i, P_i) \leq 32\beta k$.

Case 2: $\mathcal{S}_i$ is not empty. We need two steps. First we assign all pixels in each superpixel of $\mathcal{S}_i$ to R_i. Note that $\mathcal{S}_i$ is not necessarily connected, as can be seen in Fig. 3 (left). Nonetheless we can connect the superpixels in the second step using Lemma 2 below.

Lemma 2. *Let S_1 and S_2 be two superpixels in different connected components of $\mathcal{S}_i$. Let v_1 be the center of S_1 and v_2 the center of S_2. The path consisting of pixels that either intersect or border the line segment $\overline{v_1 v_2}$ must be entirely contained in R_i, and at least at twice the unit distance from the border of R_i.*

Proof. The line segment between v_1 and v_2 is contained within R_i by convexity. Similarly, the line segment from any vertex of S_1 to a vertex of S_2 is contained in R_i and necessarily also in the bounded slab that bounds these sixteen edges. Such a slab is at least as wide as S_1 and S_2 (hence it is at least $16\beta k$ pixels wide).

The line segment between v_1 and v_2 forms the spine of this slab, any pixel that intersects or borders this spine has at most two unit distance to this spine and hence is contained within the slab and via transitivity in R_i. Moreover, since the slab is at least $16\beta k$ wide, and since each pixel has distance at most two from the spine, each pixel in the path is at much more than distance two from the border of the slab and via transitivity the border of R_i. $\square$

Let S_1 and S_2 be two superpixels in different connected components of the superpixels contained in R_i. We connect S_1 and S_2 with a path of pixels according to Lemma 2. Since this path is entirely contained in R_i and since there are at least two pixels between a pixel in this path and the border of R_i, no other region will attempt to color the pixels in this path. We repeat this process until for each region the assigned pixels form a connected grid polygon and whenever we enclose an area between superpixels with these paths, we make sure to assign all the pixels in this area to R_i; by the convexity of R_i all these pixels are contained in R_i. This provides our pixel assignment P_i.

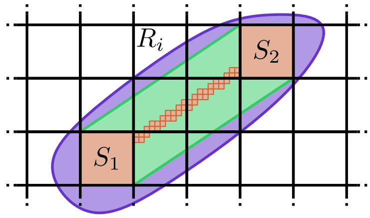

Fig. 3. A convex β-fat region R_i (purple), and the region formed by sweeping a superpixel from S_1 to S_2 (green). P_i (red) consists of S_1, S_2, and all pixels on the segment between the centers of S_1 and S_2. (Color figure online)

What remains to be proven, is that for each region R_i with non-empty $\mathcal{S}_i$, $H'(R_i, P_i) \leq 32\beta k$ holds. First, we prove that for each (boundary) point p of P_i, there is a (boundary) point q of R_i within distance $32\beta k$. By construction, we know $P_i \subseteq R_i$, so the claim holds for interior points. Now, let $p \in \partial P_i$. We assume for the sake of contradiction that there is no point of ∂R_i within distance $\sqrt{2}k$. As p is contained within R_i, we have that R_i contains the superpixels containing p, a contradiction. Second, we prove the inverse. For a point q of R_i, Lemma 1 guarantees that R_i contains a superpixel S within distance $16\beta k$ of q. Then $S \subseteq P_i$ holds, proving the claim. As $P_i \subseteq R_i$, this also proves that for each boundary point q of R_i, there is a boundary point p of P_i within distance $16\beta k$.

Theorem 2. *If $\mathcal{R}$ is a set of m β-fat convex regions for a constant β, a valid assignment exists such that for each region $R_i \in \mathcal{R}$ with a corresponding region P_i, we have $H'(R_i, P_i) = O(\sqrt{m})$. Furthermore, for any $\beta \geq 1$, there exists a set $\mathcal{R}$ of m β-fat regions such that for every valid assignment $H(R_i, P_i) = \Omega(\sqrt{m})$.*

4 Input Regions are Convex Regions

When $\mathcal{R}$ is a set of convex regions, we can easily show that the coloring has a lower-bound Hausdorff distance of $\Omega(m)$: we can place m horizontal line segments of length $\Omega(m)$ that all pass through the same pixels. Then $\mathcal{P}$ must have its elements on disjoint lines of pixels, giving Hausdorff distance at least $\Omega(m)$ for the outer regions. Each P_i must extend sufficiently far left and right. Since all P_i are connected, they will intersect a common vertical line. The topmost or bottommost intersection with this line belongs to a grid polygon with Hausdorff distance $\Omega(m)$. (Note that if the P_i need not be connected, $O(\sqrt{m})$ Hausdorff distance can always be realized.)

We will describe an algorithm that, given a set of convex regions $\mathcal{R}$, gives a set of disjoint orthoconvex grid polygons $\mathcal{P}$ such that, for all i, $H'(R_i, P_i) = O(m)$.

Observation 3. *Let $R_1, R_2 \in \mathcal{R}$ be two disjoint convex regions, and let ℓ be a horizontal line that intersects R_1 left of R_2. Then any horizontal line intersecting both R_1 and R_2 intersects R_1 left of R_2. Similarly, all vertical lines that intersect both R_1 and R_2 do so in the same order.*

Observation 3 allows us to define two partial orders $\preceq_x$ and $\preceq_y$ on $\mathcal{R}$: $R_i \preceq_x R_j$ if and only if there is a horizontal line intersecting both regions and R_i intersects the line left of R_j; since the regions are convex we get a partial order [8]. We extend this partial order as follows: first we add transitive arrows, where we recursively add the inequality $R_i \preceq_x R_j$ if there exists a region R_k with $R_i \preceq_x R_k$ and $R_k \preceq_x R_j$ and we denote this partial order by $\Pi_x(\mathcal{R})$. We then transform $\Pi_x(\mathcal{R})$ into a linear order $X_{\mathcal{R}} : \mathcal{R} \to [1, m]$ in any manner. A linear order $Y_{\mathcal{R}} : \mathcal{R} \to [1, m]$ is defined symmetrically.

Given $X_{\mathcal{R}}$ and $Y_{\mathcal{R}}$, we assign a coloring. Let Γ_k be a coarsening of Γ with $k = 2m$. For any superpixel $S \in \Gamma_k$, we denote by $S[x, y]$ the pixel that is the $(2x)^{\text{th}}$ from the left and $(2y)^{\text{th}}$ from the bottom within S. Additionally the horizontal and vertical lines induced by Γ_k are called *major lines*. Each region R_i that intersects at most one major horizontal line and at most one major vertical line is a *small region*. Each region R_i that intersects at least two major horizontal lines or at least two major vertical major lines is a *large region*. Our assignment of regions to pixels, illustrated in Fig. 4, is:

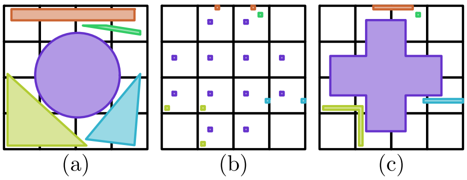

Fig. 4. The coloring algorithm for convex regions. (a) The input of four convex regions, overlaid onto a superpixel grid with $k = 10$. (b) The pixels colored in Step 1 and 2 of the algorithm. (c) The final coloring obtained after Steps 3 and 4.

1. For each small region R_i we choose one superpixel S containing a point of R_i and color the pixel $p(S, R_i) = S[X_{\mathcal{R}}(R_i), Y_{\mathcal{R}}(R_i)]$ with c_i (this single pixel will be P_i).
2. For each superpixel S and each large region R_i intersecting S that also intersects the two major horizontal *lines* incident to S, or the two major vertical *lines* incident to S, we color $p(S, R_i) = S[X_{\mathcal{R}}(R_i), Y_{\mathcal{R}}(R_i)]$ with c_i. Note that region R_i need not intersect two opposite *edges* of S.
3. For any two pixels that are colored with c_i in edge-adjacent superpixels (R_i must be large), we color all pixels in the row or column between them with c_i.
4. For any four superpixels that share a common vertex, if they each contain a pixel colored with c_i in Step 1, we color all pixels in the square between these pixels with c_i.

Let $\mathcal{P}$ be the set of polygons induced by this grid coloring.

Lemma 3. *Each polygon $P_i \in \mathcal{P}$ is simply connected.*

Proof. If R_i is small, P_i is a single pixel and thus simply connected. If R_i is large, it intersects a connected set of superpixels, and our algorithm connects all of these together, so P_i is connected. The resulting grid polygon P_i cannot contain holes: the presence of a hole would imply that the set of superpixels intersected by R_i contains a hole, which is not possible due to R_i being simply connected and convex.

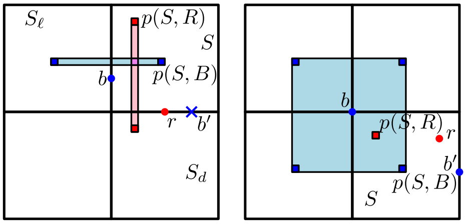

Fig. 5. The cases for the proof of Lemma 4.

Our algorithm actually produces orthoconvex polygons (refer to the full version for details).

Lemma 4. *The polygons in $\mathcal{P}$ are pairwise disjoint.*

Proof. Assume by contradiction that the colorings of two regions R and B intersect. Then the intersection was created during one of the four coloring steps. In steps 1 and 2, we assign each color to single pixels per superpixel in unique rows and columns, hence they cannot create two colorings that intersect.

Let the colorings of R and B intersect after step 3. This implies that R and B are both large regions. The intersection occurs between a vertical and horizontal pixel sequence in a super pixel S. Assume without loss of generality that the vertical sequence belongs to R and the horizontal sequence belongs to B. Consider the case that the pixel $p(S,R)$ assigned to R in S in step 2 is to the top-left of $p(S,B)$ (See Fig. 5); the other three cases are symmetric. Then the intersection occurs between the column sequence connecting $p(S,R)$ to $p(S_d,R)$ and the row sequence connecting $p(S,B)$ to $p(S_\ell,B)$, where S_d is the superpixel directly below S and S_ℓ is the superpixel directly to the left of S.

Since B is large and assigned a pixel in S it intersects both horizontal major lines incident to S or both vertical major lines incident to S. The same applies for R. We first consider the case where B does not intersect the major line through the bottom edge of S, and hence it must intersect both vertical lines. That is, B spans the vertical slab defined by S and does so in or above S. Since R intersects the cell S_d below S it then follows that $R \preceq_y B$. However, since

$p(S, R)$ lies above $p(S, R)$ we also have $B \preceq_y R$. Since $B \neq R$ we thus obtain a contradiction.

Thus, B intersects the horizontal major line ℓ through the bottom edge of S. Since R is convex, and intersects both S and S_d it intersects the bottom edge of S (and thus ℓ) in a point r. Symmetrically, B intersects the left edge of S in a point b. If B also intersects the horizontal line ℓ in some point b' this point cannot be left of r, as this would immediately imply that $B \preceq_x R$, contradicting the assignment of $p(S, R)$ and $p(S, B)$. So b' lies right of r. However, then the vertical ray starting at r pointing upwards intersects the segment connecting b and b'. Since B is convex, this segment is contained in B. This implies $R \preceq_y B$, which again contradicts the assignment of $p(S, R)$ and $p(S, B)$. It follows that step 3 does not create intersecting colorings.

Finally, let (the colorings of) R and B intersect only after step 4. Without loss of generality, the coloring of a region R is entirely contained in the coloring of a large region B. Let S be the superpixel containing the lone pixel of R. Without loss of generality we assume that the pixel $p(S, R)$ assigned to R in S is to the top-left of $p(S, B)$. Thus, B intersects S, the superpixel above S, the superpixel left of S, and the superpixel left and above S. The point b where these four superpixels meet lies inside B by convexity. Let r be any point in $R \cap S$.

As B is a large region it needs to intersect two opposite major lines incident to S. Assume that B intersects the vertical major lines, in particular the one incident to the right edge of S in a point b'. The vertical line through r intersects the segment between b and b'. The point r is above that segment, because the opposite would imply $R \preceq_y B$. As a consequence r is also right of the segment between b and b', which implies that the horizontal line through r intersects this segment left of R, a contradiction. The case where B intersects the major horizontal line through the bottom edge of S is symmetric. $\square$

If a region R_i intersects a superpixel S, then P_i has a pixel in S or in at least one of the eight adjacent superpixels. Conversely, if P_i contains a pixel in S, we know that R_i intersects S. This gives a bound on the Hausdorff distance between the regions and the grid polygons. For the boundaries, note that if R_i contains a superpixel S and all four edge-adjacent superpixels, then P_i contains S. Furthermore, if P_i contains a superpixel S, then R_i also contains S. Together this gives a bound on the Hausdorff distance between the boundaries. Since superpixels have size $\Theta(m)$, the Hausdorff distance between R_i and P_i and between their boundaries is at most $O(m)$. We thus obtain the following result.

Theorem 4. *If $\mathcal{R}$ is a set of m convex regions, a valid assignment exists such that for each region $R_i \in \mathcal{R}$ with a corresponding region P_i, we have $H'(R_i, P_i) = O(m)$. Furthermore, there exists a set $\mathcal{R}$ of m convex regions such that for every valid assignment, there exists some $1 \leq i \leq m$ with $H(R_i, P_i) = \Omega(m)$.*

5 Input Regions are General Regions

When the input regions are arbitrary, we see a sharp contrast between the case $m \leq 2$, where constant Hausdorff distance can be realized, and the case $m \geq 3$, where the Hausdorff distance may be unbounded. The fact that a single region can be represented as a grid polygon with constant Hausdorff distance was shown before by Bouts et al. [2]. In Sect. 5.1 we show that the same result holds for two regions. In Sect. 5.2 we show that for three regions, no bounded Hausdorff distance bound exists that applies to all inputs.

5.1 Two Regions

Our result for two arbitrary regions is based on a combination of two previous results: mapping a polygon to the grid with constant Hausdorff distance by Bouts et al. [2], and a result on the Painter's Problem in [15]. We briefly explain the former result in our framework using superpixels first (see Fig. 6), and then extend it to our case with two regions using the latter result.

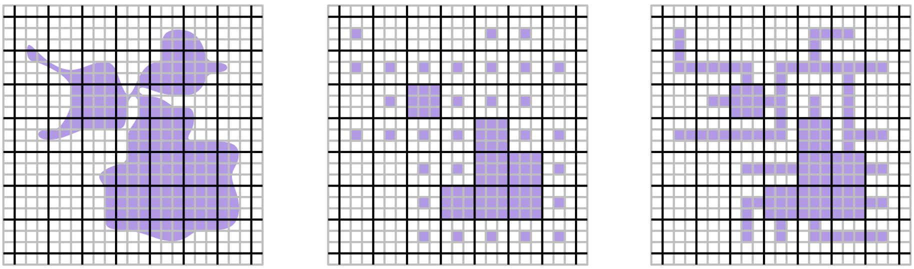

Fig. 6. Left, a region with Γ and Γ_3. Middle, the set P' of pixels chosen in the first selection. Right, the set P of pixels chosen after the spanning tree pixels are added.

Assume we have a region R that we want to represent by a grid polygon P. Consider the grid coarsening Γ_3, which has superpixels of 3×3 pixels. For every superpixel fully covered by R, choose all nine pixels in P. For every superpixel visited but not covered by R, take the middle pixel. Take nothing from superpixels not visited by R. Let the chosen pixels be P'.

Observe that P' forms a set of grid polygons that has no interior boundary cycles. Also observe that all superpixels for which at least one pixel is in P' is a connected (but not necessarily simply connected) part of Γ_3.

We make P' into one simply connected grid polygon P by using a (minimum) spanning tree on the components of P'. We will add pixels from visited superpixels only, and only ones adjacent to the already chosen center pixel. Two separate components will always be connected using one or two pixels.

Since the boundary of P does not intersect the interior of fully covered superpixels and visited superpixels always have a piece of boundary of P, it is easy to see that $H(R_i, P_i) = \Theta(1)$ and $H(\partial R_i, \partial P_i) = \Theta(1)$. This result is an alternative to the one by Bouts et al., albeit with worse constants.

A Painter's Problem instance takes a grid, and for each cell, the color white, blue, red, or purple. White indicates the absence of red and blue while purple indicates the presence of both red and blue. The question is whether two disjoint simply connected regions for red and blue exist that are consistent with all specifications of the cells, or, in the terminology of [15], "admits a painting". Since red cells can simply be colored red and blue cells blue, the problem boils down to recoloring the purple cells with red and blue pieces. The red and blue pieces in a cell provide a panel, and all panels together make up a painting. They prove:

Lemma 5 (Theorem 2 in [15]). *If a partially 2-colored grid admits a painting, then it admits a 5-painting.*

In a 5-painting each cell contains at most 5 components. The components make sure that the overall red and blue parts are connected across the whole painting. Additionally [15] show that each cell has at most 3 intervals of alternating red and blue along each side. This implies that there are only a constant number of configurations within a cell, so all configurations can be represented using a grid of constant size c for each cell.

In our problem, we have two regions R_1 and R_2 that we call red and blue, for consistency. We create a grid coarsening $\varGamma_{c+2}$. We record for every superpixel whether it is fully covered by red or blue, or visited by red and/or blue. If one color covers a superpixel completely, we assign all of its pixels to that color. If a color, say, red, visits a superpixel but blue does not, we start by making the middle $c \times c$ pixels of that superpixel red. Finally, for all superpixels visited by both red and blue, we apply the results from [15]. Since the recording of colors with panels comes from disjoint simply connected regions, namely, our input, we know that the 2-colored grid of superpixels admits a painting with connected regions/colors, so it admits one as specified in Lemma 5.

Once we choose a coloring of pixels in each 2-colored superpixel according to the panels, it remains to make the red set and blue set of pixels simply connected. The method from [15] did not produce any cycles in the 2-colored superpixels, the visited 1-colored superpixels are separate connected components of $c \times c$ pixels in the middle, and the covered 1-colored superpixels cannot create cycles either. We create a single red component by making a spanning tree of the red components. To achieve this, we only need to use pixels in the outer ring of the visited 1-colored superpixels. Then we do the same with blue. Since we add pixels of the same color to 1-colored superpixels, we will never try to color a pixel in both colors or create crossings. We then obtain the following result:

Theorem 5. *If $\mathcal{R}$ consists of two disjoint, simply connected regions, a valid assignment exists such that for each region $R_i \in \mathcal{R}$ with corresponding P_i, we have $H'(R_i, P_i) = \Theta(1)$.*

5.2 Three or More Regions

In the following we argue that the Hausdorff distance between an input of at least three general regions and any corresponding grid polygons is unbounded.

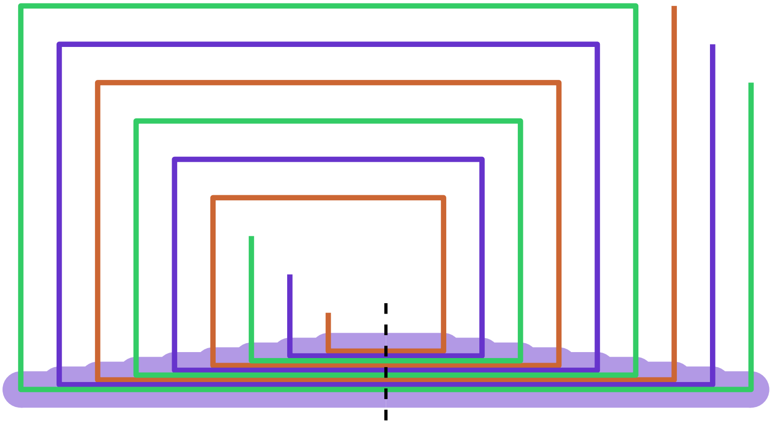

Fig. 7. The regions for $h = 3$; $\mathcal{I}$ is highlighted. The dashed line subdivides the boundary of $\mathcal{I}$ into its left and right part.

Formally, for a given integer $h > 0$, we show a construction of regions $\mathcal{R} = \{R, B, G\}$ for which there are no corresponding grid polygons with Hausdorff distance smaller than h. We only sketch the main idea here, see the full version for details.

We construct regions $\mathcal{R} = \{R, B, G\}$ that form nested spirals with a long bottleneck of height 1. The bottleneck is traversed from left to right h times by each of R, B, and G. If we remove the parts of R, B, and G inside the bottleneck, we get $3h + 3$ connected components in total. This is illustrated in Fig. 7 for $h = 3$. Outside the horizontal strip of height 1 containing the bottleneck, the three regions are more than $2h$ apart. We define the part of the plane within distance h of at least one of the bottom horizontal segments of the regions $\mathcal{R}$ as $\mathcal{I}$. All region components must be connected inside $\mathcal{I}$. Inside $\mathcal{I}$, it is possible that the grid polygons make different connections than those in $\mathcal{R}$. However, we argue that no matter how these connections are made, the grid polygons P_R, P_B, and P_G, together have to pass through $\mathcal{I}$ from left to right at least $h + 2$ times, thus requiring $\mathcal{I}$ to have height at least $2h + 3$. However, the available vertical space is only $2h + 1$ if the Hausdorff distance must stay below h, allowing $h + 1$ connections of pixel polygons. Hence, we obtain a contradiction.

The most involved part is to argue that P_R, P_B, and P_G, together have to pass through $\mathcal{I}$ at least $h + 2$ times. This argument critically depends on the following Lemma (see Fig. 8 for an illustration).

Lemma 6. *Given an alternating sequence* $V = r_1, b_1, g_1, ..., r_k, b_k, g_k$ *of* $3k$ *3-colored points on a line, any planar drawing below the line connecting points of the same color induces a partition of the points into at least* $2k + 1$ *components.*

The idea is that $\mathcal{I}$ splits the regions in $\mathcal{R}$ (and thus their corresponding grid polygons) into $3h + 3$ connected components. However, the regions intersect the

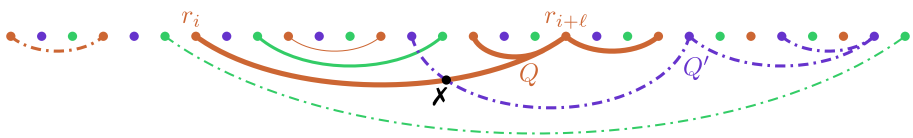

Fig. 8. A set Q that includes two red points r_i and $r_{i+\ell}$ splits V into two disjoint subsequences V_1 and V_2, that have at most one set, namely Q, in common. If there was a second such a set Q', the grid polygons corresponding to Q and Q' would intersect. (Color figure online)

right half of the boundary of $\mathcal{I}$ only $3h$ times, and in an order in which the colors alternate, we can use Lemma 6 to show that we can decrease the number of connected components by at most $h-1$ by connecting the regions incident to "the right side" of $\mathcal{I}$ to other regions on the right side of $\mathcal{I}$. The same holds for the regions on the left side of $\mathcal{I}$. It thus follows that the remaining $3h-2(h-1) = h+2$ of the reduction in the number of connected components (after all, in the end there are only three regions left) must be achieved by connecting regions incident to "the left side of" $\mathcal{I}$ to "the right side" of $\mathcal{I}$. Therefore, P_R, P_B, and P_G pass through $\mathcal{I}$ at least $h+2$ times as claimed. Therefore, this allows us to obtain the following result:

Theorem 6. *For all integer $h > 0$ there exist three regions $\mathcal{R} = \{R_1, R_2, R_3\}$, for which there is no valid assignment to grid polygons P_1, P_2, P_3 so that all regions $R_i \in \mathcal{R}$ have $H(R_i, P_i) < h$.*

6 Conclusion

In this paper we have shown what Hausdorff distance bounds can be attained when mapping disjoint simply connected regions to the unit grid. We expressed our bounds in terms of the number of regions and obtained different results depending on the shape and size characteristics of the regions, and showed that they are worst-case optimal. The result in Sect. 5.1 generalizes a result of Bouts et al. [2] and the result in Sect. 5.2 shows that a result by Van Goethem et al. [15] cannot be generalized from two to three colors. Our results are slightly more general than we expressed them: for example, the bound for point regions in fact holds for any set of regions that each have constant diameter.

We assumed that our regions all had the same shape and size characteristics. In some cases it is interesting to see what happens in combinations. In particular, suppose we have one general region R_0 and m point regions $R_1, \ldots, R_m$; what Hausdorff bounds can be attained? It turns out that we get a trade-off: we can realize a Hausdorff distance of $O(\sqrt{m})$ for the point regions and for R_0, but we can also realize a Hausdorff distance of $O(1)$ for R_0 but then some point region will have a Hausdorff distance of $\Theta(m)$. Figure 9 illustrates this. We may map the points to the grid first using the $O(\sqrt{m})$ bound, and then map R_0, or we can map the points to the grid in a constant width strip close to the boundary

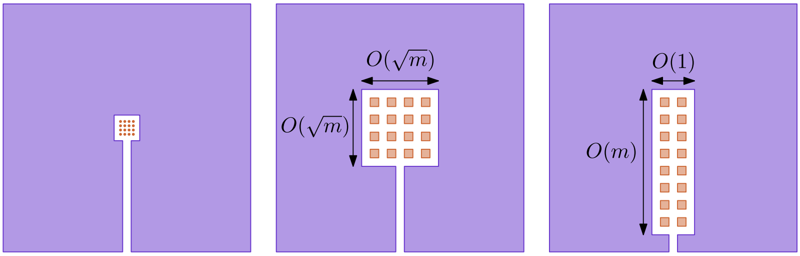

Fig. 9. Left, an instance with one general region (purple) and m point regions. Middle and right, two possible realizations for different Hausdorff bounds.

of R_0. Note that in the former case, we could have left a spacing of three pixels between the mappings of the point regions. Then the point regions still attain the $O(\sqrt{m})$ bound, while $H(R_0, P_0) = O(1)$ by using the extra space to allow P_0 to reach every necessary place. However, $H(\partial R_0, \partial P_0)$ will still be $\Theta(\sqrt{m})$, so we do not improve $H'(R_0, P_0)$.

While we concentrated on worst-case optimal bounds, our constructive proofs of the upper bounds will often give visually unfortunate output. Also, for a given instance we may not achieve $O(1)$ Hausdorff distance for m point, β-fat convex, or convex regions even when constant would be possible for that instance. This leads to the following two open problems. Firstly, can we realize visually reasonable output when this is possible for an instance (and how do we define this)? Secondly, can we realize a Hausdorff distance that is at most a constant factor worse than the best possible for each instance, in polynomial time?

Acknowledgments. This research was partially supported by the Dutch Research Council under the projects 612.001.651, 614.001.504 and 628.011.005.

References

1. Althaus, E., Eisenbrand, F., Funke, S., Mehlhorn, K.: Point containment in the integer hull of a polyhedron. In: Proceedings 15th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 929–933 (2004)
2. Bouts, Q.W., Kostitsyna, I., van Kreveld, M., Meulemans, W., Sonke, W., Verbeek, K.: Mapping polygons to the grid with small Hausdorff and Fréchet distance. In: Proceedings 24th Annual European Symposium on Algorithms, pp. 1–16 (2016)
3. Christ, T., Pálvölgyi, D., Stojaković, M.: Consistent digital line segments. Discret. Comput. Geom. **47**(4), 691–710 (2012)
4. Chun, J., Kikuchi, K., Tokuyama, T.: Consistent digital curved rays. In: Abstracts 34th European Workshop on Computational Geometry (2019)
5. Chun, J., Korman, M., Nöllenburg, M., Tokuyama, T.: Consistent digital rays. Discret. Comput. Geom. **42**(3), 359–378 (2009)
6. de Berg, M., Halperin, D., Overmars, M.: An intersection-sensitive algorithm for snap rounding. Comput. Geom. **36**(3), 159–165 (2007)

7. Goodrich, M.T., Guibas, L.J., Hershberger, J., Tanenbaum, P.J.: Snap rounding line segments efficiently in two and three dimensions. In: Proceedings 13th Annual Symposium on Computational Geometry, pp. 284–293 (1997)
8. Guibas, L.J., Yao, F.F.: On translating a set of rectangles. In: Proceedings 12th Annual ACM Symposium on Theory of Computing, pp. 154–160 (1980)
9. Harvey, W.: Computing two-dimensional integer hulls. SIAM J. Comput. **28**(6), 2285–2299 (1999)
10. Hershberger, J.: Stable snap rounding. Comput. Geom. **46**(4), 403–416 (2013)
11. Klette, R., Rosenfeld, A.: Digital Geometry: Geometric methods for digital picture analysis. Elsevier (2004)
12. Klette, R., Rosenfeld, A.: Digital straightness - a review. Discret. Appl. Math. **139**(1–3), 197–230 (2004)
13. Löffler, M., Meulemans, W.: Discretized approaches to schematization. In: Proceedings 29th Canadian Conference on Computational Geometry (2017)
14. Löffler, M., Simons, J.A., Strash, D.: Dynamic planar point location with sub-logarithmic local updates. In: Proceedings 13th International Symposium on Algorithms and Data Structures, pp. 499–511 (2013)
15. van Goethem, A., Kostitsyna, I., van Kreveld, M., Meulemans, W., Sondag, M., Wulms, J.: The painter's problem: covering a grid with colored connected polygons. In: Frati, F., Ma, K.-L. (eds.) GD 2017. LNCS, vol. 10692, pp. 492–505. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-73915-1_38