

Drawing Depth Contours with Graphics Hardware

Ian Fischer
Harvard University

Craig Gotsman
Technion – Israel Institute of Technology

Abstract

Depth contours are a well-known technique for visualizing the distribution of multidimensional point data sets. We present an image-space algorithm for drawing the depth contours of a set of planar points. The algorithm is an improvement on existing algorithms based on the duality principle, implemented using 3D graphics rendering techniques.

1 Introduction

Given a set P of n points in R^d , the *location depth* of a point $q \in R^d$ relative to P describes, intuitively, the relationship of q to the distribution of the points in P . The associated concept of *depth contour* is the locus of all points with the same fixed location depth. Drawing nested depth contours helps to visualize the distribution of the points. In particular, the *Tukey median* is the centroid of the deepest such contour. See example in Figure 1. Location depth and depth contours have applications in robust statistics, hypothesis testing, and some areas of cell biology. In this paper we will deal with the planar case, i.e. $d=2$. We start with some formal definitions.

Definitions: Let $P = \{p_1, p_2, \dots, p_n\}$ be a finite set of points in R^2 and q be an arbitrary point in R^2 . The *location depth* of q relative to P , denoted by $ld_P(q)$, is the minimum number of points of P lying in any closed half plane determined by a line through q . The *k -th depth contour* of P , denoted by $dc_P(k)$, is the boundary of the set of all points q in R^2 with $ld_P(q) \geq k$. \diamond

The quantity $ld_P(q)$ is an integer in the range $\{0, \dots, n\}$. It is 0 when q lies outside the convex hull of P and n when all points coincide with q . The k -depth contour is sometimes called the *k -hull*, and points on $dc_P(k)$ can be shown to be *k -splitters* of P , because every line through q has at least k points of P lying on or above it and at least k points of P lying on or below it.

It is easy to verify a number of simple properties of depth contours: $dc_P(k)$ is a convex polygon for any k , $dc_P(1)$ is the boundary of the convex hull of P , and $dc_P(k+1) \subseteq dc_P(k)$. It is less obvious to show that there will always be a depth contour of depth $\lfloor n/3 \rfloor$ but not necessarily one of depth $\lfloor n/2 \rfloor$.

2 Computing Depth Contours

2.1 The geometric algorithm

Miller *et al* [4], improving results of Cole *et al* [2], present an algorithm for computing the depth contours and related entities for a set of points. The algorithm itself is fairly simple, makes extensive use of duality, and proceeds as follows: Given a set P of points, map them to their dual arrangement of lines. Then apply topological sweep to find the planar graph of the arrangement and label the vertices of the arrangement with their *levels* – the number of dual lines above them. The *depth* of a vertex is $\min(\text{level}(v), n - \text{level}(v) + 1)$. Finally, for a given k , compute $dc_P(k)$ by finding the lower and upper convex hull of the vertices at depth k . Each such vertex corresponds to a half-plane in the primal plane, and $dc_P(k)$ is the boundary of the intersection of these half-planes (which might be empty, in which case $dc_P(k)$ does not exist).

The complexity of this algorithm is $O(n^2)$ time and $O(n^2)$ space, which has been shown to be optimal. A related algorithm for computing the location depth of a single query point in $O(\log n)$ time follows, and this is also optimal. The so-called *bag plot* of the points, which is the convex region containing no more than half the points with largest depth, may also be computed in $O(n)$ time once the depth contours have been computed. Finally, the Tukey median can be easily computed once the deepest depth contour is known, or independently in optimal $O(n \log n)$ expected time using the randomized algorithm of Chan [1].

2.2 The image-space algorithm

For large n , an $O(n^2)$ time geometric algorithm could be prohibitive, especially in an interactive application where the point set is dynamic. To address this, Krishnan *et al* [3] present an image space algorithm to approximate all the depth contours of a given point set. Essentially, it takes advantage of the graphics hardware to *draw* an image of the depth contours as a set of colored pixels. Thanks to this, their algorithm outperforms the geometric algorithm by at least one order of magnitude. But this efficiency comes at the price of losing sub-pixel information. However, as we shall see later, this error is minimal and the results are still quite acceptable.

The image-space algorithm is based on the same duality principle as the geometric algorithm. It consists of two phases. In the first phase, the input point set P is scan-converted to lines in the dual image plane. Since the dual plane is discrete, it is possible to compute the *level* of each pixel. This information is used in the second phase to create the depth contours.

In more detail, the algorithm proceeds as follows. Without loss of generality, assume that each point in P is located inside the square of edge length two centered at the origin. In the first phase, transform each point of P to its line in the dual plane, and define the *level* of a point q in the dual plane to be the number of lines below or passing through it, and the *depth* of q to be $\min\{n - \text{level}(q), \text{level}(q)\}$. Because our dual plane is of finite size, it is necessary to guarantee that all intersection points of the lines lie in this finite region. Krishnan *et al* claim that by using two separate bounded dual planes of size 2×4 – bounded dual 1 (BD1) maps (p_x, p_y) to $y = -p_x x + p_y$ and bounded dual 2 (BD2) maps (p_x, p_y) to $y = -p_y x + p_x$ – all intersections must lie in one of these two dual planes. Thus, in practice, the algorithm runs on both bounded duals. The level of each point in the dual plane is computed by drawing each dual line in turn and incrementing the region above the dual line by one in the stencil buffer. This is accomplished by drawing the entire half-plane lying above the dual line, which causes all the pixels on or above the line to be incremented by one if the stencil operation is set to increment. The dual lines themselves are drawn into a separate buffer in order to keep track of which pixels actually contain lines.

Phase two uses the knowledge of the level of each point in the dual plane to compute the depth contours. The image of the dual lines is scanned, and for each pixel q on a dual line, the corresponding primal line at fixed z -depth $\min\{n - \text{level}(q), \text{level}(q)\}$ is rendered as a colored 3D graphics primitive using the z -buffer. Since n is fixed and the value $\text{level}(q)$ is available in the stencil buffer, the appropriate depth is easily determined. The hardware depth test is set to LESS while rendering these primal lines, and the line color should be distinct for each depth. The resulting rendered image will contain the depth contours of the point set P as the boundaries between colored regions.

The algorithm generates a discrete image of the depth contours and thus aliasing is possible. The three main sources of error are: (1) sub-pixel precision in the input (i.e. the coordinates of the input points do not necessarily coincide with the pixel grid points), (2) computing the depth of a pixel in the dual plane, and (3) rendering the lines back in the primal. Krishnan *et al* prove that the depth contours produced err by at most one pixel. Such a small error is acceptable for most applications

and the runtime speedup is usually worth the small discrepancy.

The complexity of this algorithm is $O(nm + nm^{1/2} + m^{3/2}) = O(nm + m^{3/2})$, where m is the number of pixels in the output image. The first term corresponds to the first stage of the algorithm – n half-planes are rendered, each of which covers $O(m)$ pixels. The second term corresponds to the second stage of the algorithm, where n lines are rendered, and each of them has a length of $O(m^{1/2})$. The final term also corresponds to the second stage of the algorithm – there are a maximum of m pixels that represent primal lines that are rendered, and each of them will cover $O(m^{1/2})$ pixels when this happens.

3 Our Algorithm

Our algorithm is an improvement of the image-space algorithm of Krishnan *et al* [3]. That algorithm renders a 3D *line* in the primal image for each pixel drawn in the dual plane, and these lines cover regions in the primal image plane. However, the following simple observation results in a number of improvements: The union of primal lines corresponding to the points along a *dual line segment* between two adjacent vertices of the arrangement form a *double wedge*, which may be rendered as two *triangles* at a fixed depth. The accurate identification of these dual line segments is the only difficult part.

The improvement is threefold: First, the number of rendering operations is dramatically decreased. We render a reasonable number of triangles ($O(n^2)$) instead of a huge number of lines ($O(m)$). This improvement is particularly significant if n is much smaller than m . Second, the accuracy of the result is improved. Since we render an entire triangular region instead of covering it by lines, we are less prone to aliasing problems if the lines are not dense enough. Third, we only have to deal with $O(nm^{1/2})$ pixels from the dual plane renderings, rather than the full m pixels required by the original algorithm. This also implies that to make the best use of this algorithm, the number of sites should be no greater than $m^{1/2}$. This is reasonable, as having more sites than that results in a high density of depth contours relative to the number of pixels, meaning that the resolution of the image is too small anyway to support the required amount of detail.

3.1 Detecting dual line segments

The main challenge in implementing this algorithm is accurately detecting the dual line segments. Because the lines have been rasterized, it can be difficult to determine exactly which direction a given line is heading, or even whether there is a line segment between two neighboring pixels. We deal with the first issue in a

fairly unsophisticated manner that could be improved upon, and with the second problem by sidestepping it. We use the following algorithm to trace the line segments in each of the dual planes. First, when rendering just the lines of the bounded duals, use the same stencil buffer operation as used when rendering the half-planes. This will cause the stencil buffer to be 0 wherever no line has been drawn, 1 wherever a line has been drawn with no intersecting lines, and greater than 1 at every point of intersection between two or more lines. Clearly, this allows us to determine where intersections between lines occur, and therefore where a given line segment terminates.

Once we have the stencil buffers from the bounded dual half-planes, the stencil buffers from the bounded dual lines, and the color buffers from the bounded dual lines, we can begin tracing the bounded dual line segments and computing the appropriate set of double wedges. First, we find a pixel where a line intersects the edge of the bounded dual. We identify that pixel as one end of a line segment and push it on an event queue, and then walk along the line by finding a neighboring pixel that is “colored in” (i.e., has had a line rasterized there). As we walk along the line, we erase it so that we do not consider it again in the future, thereby guaranteeing termination of the algorithm. We continue until we encounter a pixel whose stencil buffer value is greater than 1, at which point we flag that pixel as the end of the line segment, and create the appropriate double wedge, assigning the depth found in the half-plane stencil buffer along that segment between the endpoints (the pixels at the endpoints may have a different value for the level, so we do not use them). Every time an endpoint is found, it is pushed on the event queue, and every time a double wedge is created, we check if it has any neighboring pixels that are colored (thereby signifying that there is some line segment leading from that endpoint that has not been considered). If so, we start a new double wedge from that end point. If not, we pop the next event from the queue. We repeat this until the queue is empty, at which point we look for the next pixel where a line segment intersects the edge of the bounded dual. We repeat this until there are no more line segments intersecting the edge of the bounded dual. Once we have completed this for both bounded duals, we draw the set of double wedges at the appropriate depths.

3.2 Caveats

Problems with our algorithm can arise near points of intersection – the rasterizations of two lines will intersect at some set of pixels, but there may be places near those intersection pixels where the rasterizations border each other, making it difficult to determine which direction the current line segment is heading from a given pixel. A wrong choice can easily result in an incorrect

double wedge, as the point of intersection between the two lines in question can be completely missed, resulting in a “line segment” that has a corner. We deal with this problem by considering only two or three of the eight neighboring pixels when walking along a line. Specifically, the first colored pixel we find adjacent to an endpoint reduces the set of directions we can travel from eight to three, then the next pixel we see in that direction that is either immediately clockwise or counter-clockwise from the first direction further limits the set of directions to those two directions. All other directions will be ignored. This works in most cases, but it is still possible to design arrangements of rasterized lines for which it will fail.

3.3 Complexity


Our line-tracing algorithm runs in $O(nm^{1/2})$ time – for every dual line we typically walk through $O(m^{1/2})$ pixels. There will be $O(n^2)$ double wedges to render – an arrangement of n lines has $O(n^2)$ line segments – so drawing the double wedges requires $O(n^2)$ time on the CPU. It can also be seen quite easily (by induction) that for each site, drawing the double wedges associated with that site results in exactly (up to rasterization error) m pixels being covered between all of the wedges. Thus, rasterizing the $O(n^2)$ double wedges takes $O(nm)$ time. The maximum number of primal lines we can draw due to the problem of neighboring intersection points is $O(n^2)$, since those can only be drawn from intersection points, which gives a total complexity of $O(n^2m^{1/2})$. The initial steps of the algorithm (drawing the dual half-planes and dual lines) has the same complexity in our algorithm as in that of Krishnan *et al.* – $O(nm + nm^{1/2})$. Thus, the total worst-case complexity of our algorithm is $O(nm + n^2m^{1/2})$. In the worst case that n is $O(m^{1/2})$, this simplifies to the same complexity as the algorithm of Krishnan *et al.* In the more realistic case that $n = o(m^{1/2})$, our algorithm is asymptotically faster.

4 Experimental Results

To experimentally compare our image-space algorithm with that of Krishnan *et al* [3], we implemented and ran both on a variety of test cases. Both algorithms may be implemented quite easily in C++ and OpenGL. However, a couple of important implementation details are missing from [3], worth mentioning here. First, and most important, was the transformation of a point in BD2 (the second bounded dual) back to a line in the primal plane. Without this the algorithm is meaningless, as the standard type of transformation would give incorrect lines in the primal plane. Thus a point p in BD2 reverts to the line $y = x/p_x - p_y/p_x$. The implementation of Krishnan *et al* simply rotates BD2 by 90 degrees to get the same effect. A second issue is one of coverage of the primal plane with the lines corresponding to points from the two bounded duals. Krishnan *et al* do

not discuss why, or even whether, the algorithm will result in full coverage of the primal plane, or if the lines of one depth contour will completely occlude the lines of contours deeper than it, and be completely occluded by the lines of contours shallower than it. In practice, the primal plane is completely covered after only two or three sites. However, we are not so lucky when trying to cover deeper levels, so often there are significant artifacts where a given level shows the color of the next deeper level because the level was not fully covered. See Figure 1 for an example of inadequate coverage. Ultimately, the only way to deal with this problem is to draw the lines at higher density, i.e. draw multiple lines per pixel by “supersampling” the pixel. If a pixel is colored in the dual arrangement, we subdivide it into a regular grid of s by s subpixels, each of which we consider to also be colored, thus generating s^2 primal lines at this depth. Unfortunately, a suitable value of s is dependent on both m and n . Krishnan *et al* deal with the coverage problem in their implementation by increasing the line thickness, but this seems also to be a significant hidden source of inaccuracy.

We ran the two algorithms on an Intel Centrino 1.8GHz processor with 1GB of RAM and a 128MB RAM ATI FireGL128 video card running at a resolution of 1000x1000 pixels. Our test results are the average of ten runs on different randomly generated sites at 100, 250, 500, 750, 1000, and 2500 sites. Our experiments show that our algorithm outperforms the algorithm of [3] by up to a factor of two until the number of sites rises above the square root of the number of pixels, at which point they become essentially equivalent.



5 Conclusion

Our fairly simple improvements to the algorithm of Krishnan *et al* [3] allow us to more quickly draw the depth contours and Tukey median of a set of sites in 2D. They also allow us to visualize the depth contours in more detail, as the algorithm shows an increasing relative efficiency as the resolution of the rendered image increases while the number of sites remains constant. Indeed, our algorithm is now fast enough that we can view interactive changes to the depth contours for a small set of dynamic sites.

Acknowledgement: Thanks to Vikas Goela for his help in implementing the algorithm described in this paper.

References

- [1] T. M. CHAN. An optimal randomized algorithm for maximum Tukey depth. *Proc. SODA*, 2004.
- [2] R. COLE, M. SHARIR, AND C. K. YAP. On k -hulls and related problems. *SIAM J. Comp.* 15(1):61–77, 1987.
- [3] S. KRISHNAN, N. MUSTAFA, AND S. VENKATASUBRAMANIAN. Hardware-assisted computation of depth contours. *Proc. SODA*, 2002.
- [4] K. MILLER, S. RAMASWAMI, P. ROUSSEEUW, T. SELLARES, D. SOUVAINÉ, I. STREINU AND A. STRUYF. Fast implementation of depth contours using topological sweep. *Statistics and Computing*, 13:153-162, 2003.




Figure 1: Portion of depth contour image (900x900 pixels) generated by our implementations of the algorithm of [3] (left) and our algorithm (right). White points are (a subset of the) sites, and dark point is their Tukey median. The gray point in the bottom right is the centroid of the sites. The algorithm of [3] generated 637,310 primal lines using a supersampling rate of 2, compared to 64,324 lines and 15,398 double wedges at the same supersampling rate for our algorithm. Our algorithm runs about twice as fast for this input. Note the aliasing artifacts in the left image compared to the right.