Filling Space with Random Line Segments

John Shier

Abstract. The use of a nonintersecting random search algorithm with objects having zero width ("measure zero") is explored. The line length in the units placed is taken as a scale-free negative-exponent power law versus the placement number. In the examples studied computational evidence indicates that the algorithm is nonhalting.

1. Introduction

In previous work [1] a space-filling geometric algorithm was described, based on (a) a sequence of shapes whose areas obey a negative-exponent power law (Eq. (1) below) versus shape number *i*, and (b) placement of successive ever-smaller shapes by nonoverlapping random search. Figure 1 shows a circle example.

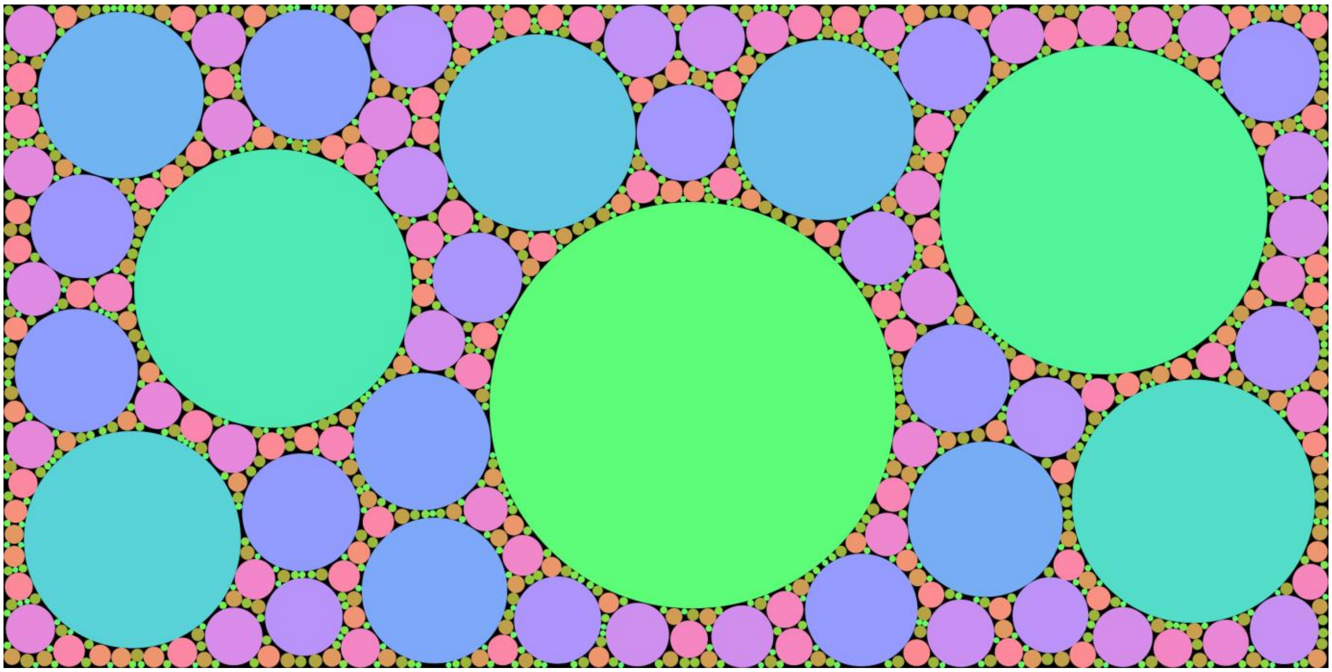

Fig. 1. A space-filling circle fractal. $c = 1.48$, $N = 3$, 1000 circles, 94% fill.

Fractals like Fig. 1 have several interesting properties. (1) The pattern is space-filling by construction, i.e., the sum of all the circle areas (to infinity) equals the area of the bounding region. (2) The probability that any circle touches any other is vanishingly small, i.e., the shapes are non-intersecting. (3) The maximum distance from any point in the unfilled region (the black area in Fig. 1) to the nearest circle or the boundary falls steadily as the algorithm moves forward, approaching a limit of 0.

The present study asks a different question: If one attempts to place nonintersecting line segments[1] within a bounding region does property (3) still hold?

Straight lines are primary elements of traditional geometry. One of their properties is that two (Euclidean) lines always have exactly one intersection point unless they are parallel. If we wish to have a nontrivial intersection test we cannot have all of the segments parallel to each other.

The *i*-th circle area for Fig. 1 is given by [1]

[1] The idea of using "measure zero" objects instead of ones with finite areas was first suggested by Barry Cipra.

$$A_i = \frac{A_b}{\zeta(c,N)(i+N)^c} \qquad (1)$$

where A_b is the area of the bounding region, c and N are parameters, and $\zeta(c,N)$ is the Hurwitz zeta function. In the present study we assume a sequence of segment lengths L_i ($i = 0, 1, \ldots$)

$$L_i = \frac{N^c}{(i+N)^c} L_0 \qquad (2)$$

where L_0 is the length of segment 0, and c and N are parameters which define the way the line segment lengths decrease with increasing i. The power-law sequence of Eq. (2) is "scale-free" or "fractal" by some definitions. If $c > 1$ the total length of all the segments is finite, and thus the segments can hardly be space-filling in the sense of property (3) above. If $c < 1$ the total length of the first i segments increases without bound as i increases. We assume $c < 1$ for this study. The nonintersecting random search algorithm of [1] is retained (Appendix A has a flow chart).

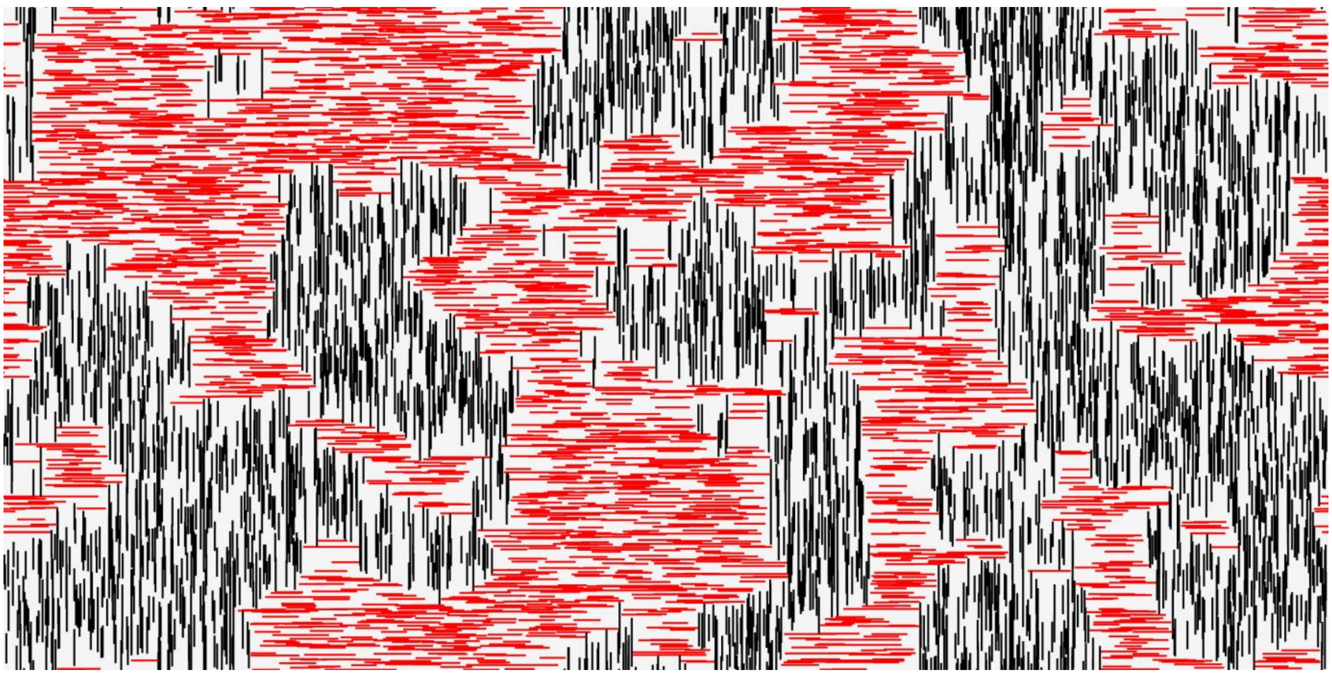

Fig. 2. A line segment fractal with the segments alternating between vertical and horizontal. The bounding region has width 12 and height 6, with periodic boundaries [1]. $c = .333$, $N = 5$, $L_0 = 1.2$, 6000 line segments. The line segments have an exaggerated width so that they can be seen more easily.

Figure 2 shows an example of the algorithm in which the line segments alternate between vertical and horizontal, with periodic boundaries such that the pattern is a tile which can be seamlessly joined to identical tiles. The bounding region is divided into two areas: "vertical" and "horizontal". There is a "coastline" between the two areas somewhat similar to real coastlines ([2] chapter II), with the coastline having bays and promontories at all length scales down to the smallest line segment length. The property that any point not contained in one of the line segments has a maximum distance to the nearest line segment appears to hold qualitatively. The largest voids appear near boundaries between the "red" and "black" areas.

It is not difficult to see what is going on here. A "vertical" line is much more likely to find a placement near other vertical lines (which it cannot intersect with) rather than near horizontal ones. The average number of random trials needed to find a nonintersecting placement is only slightly above 2. The algorithm does not halt because there is always a place for a short new line segment in an area with lines parallel to it.

2. Line Segments with Random Orientation.

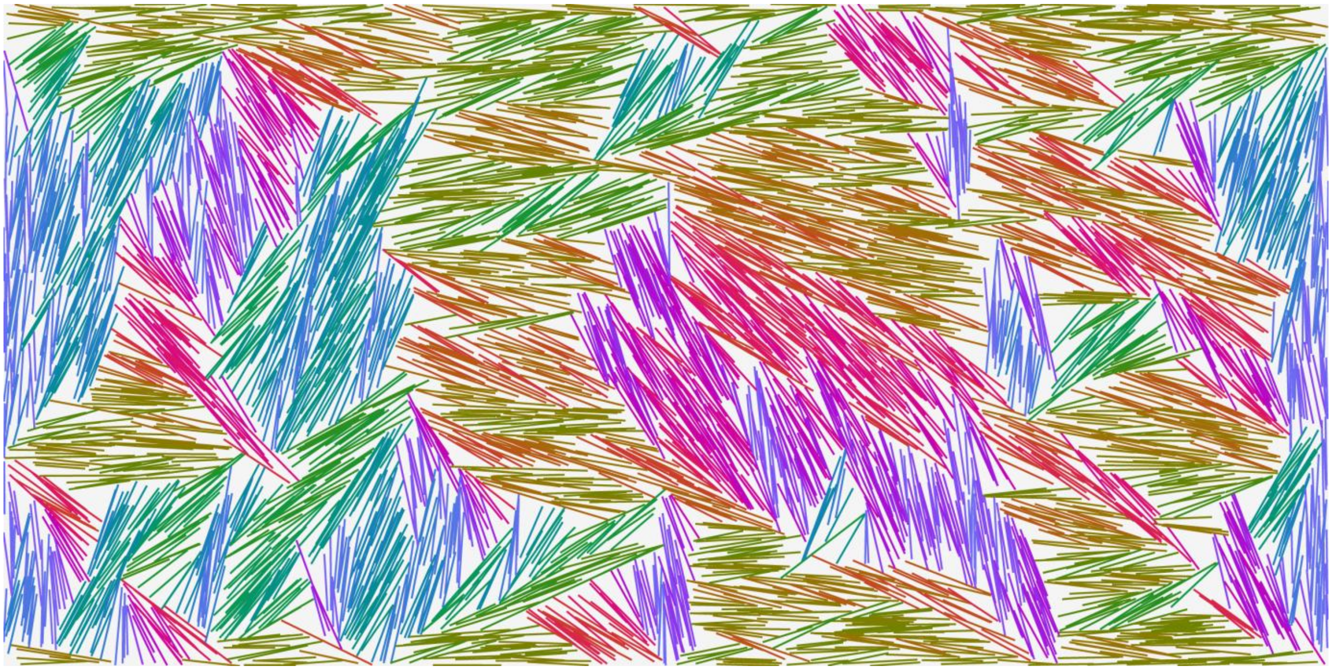

Fig. 3. A line segment fractal with the segments having random orientation angles. The bounding region has width 12 and height 6, with inclusive boundaries [1]. $c = .25$, $N = 5$, $L_0 = 1.2$, 4000 line segments. The color is continuously periodic in the orientation angle, with lines which have the same slope having the same color. 193,692 trials were needed to place the 4000 segments.

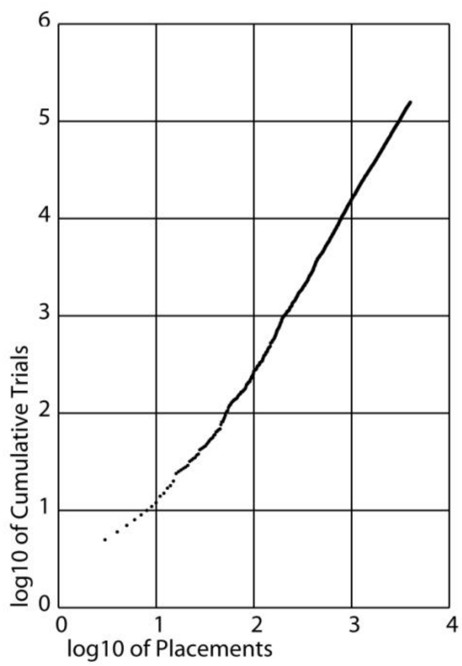

Fig. 4. Trial statistics for segments having random orientation angles. The bounding region has width 12 and height 6, with inclusive boundaries [1]. $c = .25$, $N = 5$, $L_0 = 1.2$, 4000 line segments.

In Figure 3 the line segments have not only a random *x* and *y* at each trial, but also a random orientation angle. The effect of this is that the random search algorithm must on average run longer to find a placement, and the placement will usually be in an area where the other segments have nearly the same orientation. The random placement of the first few segments establishes regions with a preferred orientation. Because an inclusive boundary is used, the line segments near the boundary are nearly parallel to it.

The halting question does not have such an obvious answer as for the orthogonal case (Fig. 2). Figure 4 shows a log-log plot of the cumulative number of trials versus placement number. As the number of segments increases the trend settles down to a straight line in log coordinates, which indicates that the relationship is a positive-exponent power law with an exponent equal to the slope of the line. This is very similar to what is seen for fractals with 2D shapes [1]. The limiting slope (i.e., the power-law exponent) here is a bit less than 2.0. Since a power law never becomes infinite, this provides computational support for the idea that the algorithm does not halt. A formal proof of this offers a serious challenge.

The blend of randomness and order seen in Figs. 2 and 3 has some visual charm, and will perhaps find a place in decorative art.

3. Tepees.

The possibilities expand when we consider placement of units which have more than one line segment. In Figure 5 the basic element is an inverted V ("tepee") made of two line segments. The effect of the algorithm is to "nest" the tepees. The right segment of each tepee is colored a modulated dark gray and the left segment a modulated light gray, which results in shading effects which can create illusions. When the image is viewed close up it may look like piles of thin rectangular sheets with square pits among them. Viewed from a greater distance it may look like random mountain ridges with light coming from the left and blunt ends to the ridges. In the ridge view, the crests form nearly continuous meandering lines.

Since there is no "real" object here, the eye-brain system constructs one. Adjustment of the parameters and the color scheme can give a variety of visual effects.

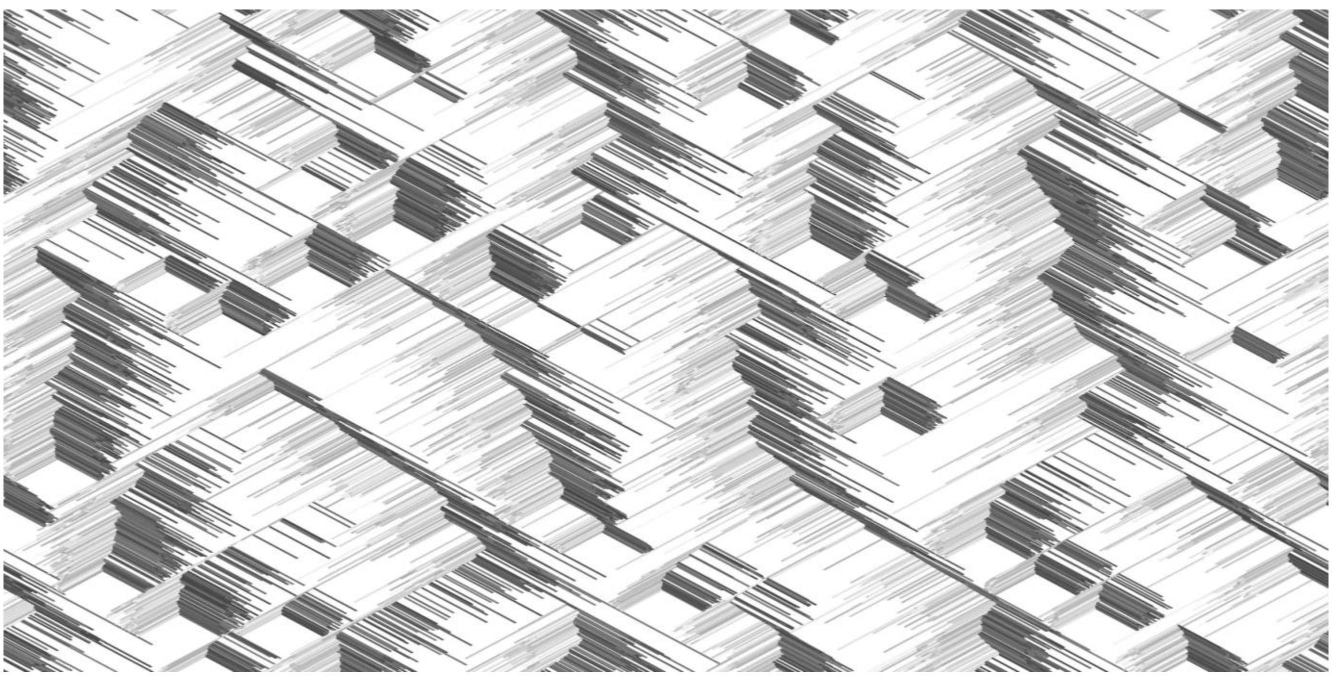

Fig. 5. A line segment fractal with inverted Vs. The bounding region has width 12 and height 6, with periodic boundaries [1]. $c = .27$, $N = 5$, $L_0 = 1.5$, 6000 line segments. Placement of the 6000 tepees required 3,707,368 trials.

4. Two Elements -- └ and ┐.

The two elements are placed alternately -- an L followed by an inverted L, etc. The result is shown in Fig. 6. As in Fig. 5 there is strong nesting, but it is selective. The new Ls are nested with previously placed Ls, and similarly for the inverted Ls. The color scheme has all of the vertical line segments in modulated dark gray, and the horizontal segments in modulated light gray. Since there is nothing real in this image, the eye-brain system constructs something. Most people see a series of rocky ledges separated by vertical scarps. With this color scheme the tops and bottoms of the scarps have strong contrast and a scarp edge seems to form a continuous noisy curve.

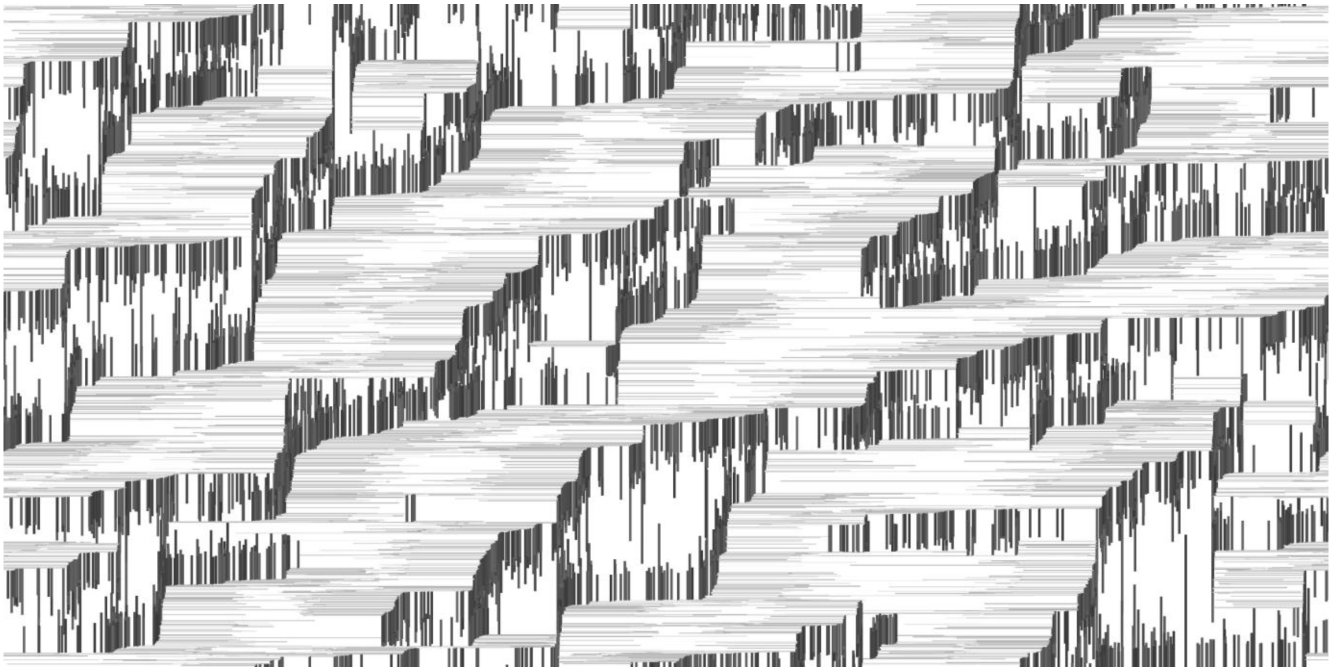

Fig. 6. A line segment fractal with alternating Ls and inverted Ls. The horizontal and vertical segment lengths have a 2:1 ratio. The bounding region has width 12 and height 6, with periodic boundaries [1]. $c = .27$, $N = 5$, $L_0 = 1.5$, 6000 line segments. Placement of the 6000 elements required 5,461,875 trials.

5. Plusses.

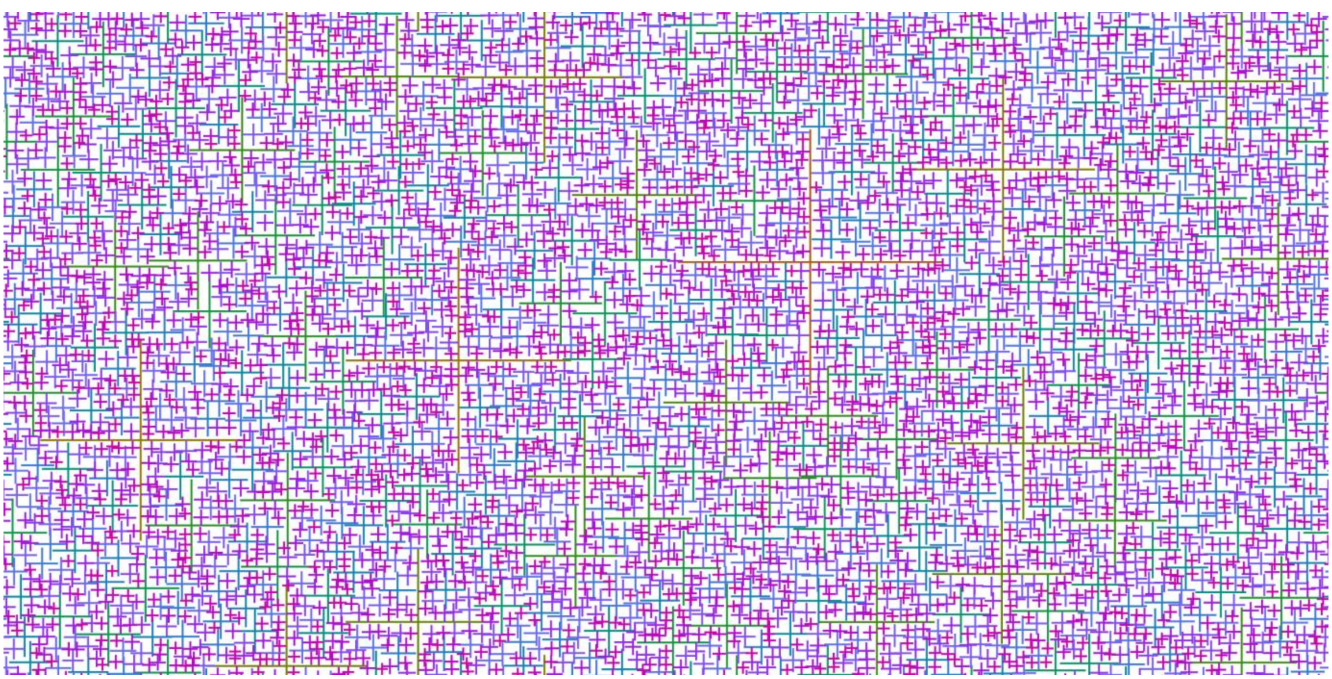

Fig. 7. A line segment fractal with plusses. The bounding region has width 12 and height 6, with periodic boundaries [1]. c = .40, N = 2, L_0 = 1.2, 5000 plusses. Log-periodic color. Placement of the 5000 plusses required 720,965 trials.

This two-line unit produces a rather different pattern (Fig. 7). With the log-periodic color scheme used elements of a given size have the same color. The most obvious thing is that there is no nesting. The two orthogonal arms of the plusses rend to repel other plusses, resulting in a pattern of ever-finer features with only very small voids. If c becomes too small the algorithm does halt. With too-low c the size of the plusses falls very slowly, and the size of the voids where they can be placed decreases too rapidly.

6. Discussion.

Familiarity with the Shier-Bourke paper [1] will help in understanding this work. As in [1], quite simple assumptions can produce complicated results.

Traditional plane geometry is much concerned with the intersection points of lines, their intersection angles, and the distances between vertices. The present study works with line segments and units composed of two line segments. It can be thought of as the geometry of nonintersecting lines, or geometry without vertices.

It is conjectured that the algorithm does not halt for any of the examples described here. Computation (Fig. 4 is an example) supports this.

A number of the statements made (like the maximum distance claims) seem obvious but have no formal proof.

These structures will have a fractal dimension D. The value(s) of D remain to be determined. Since these fractals are constructed with zero-width lines, they invite comparison with well-known space-filling fractal curves such as those of Peano and Hilbert [2]. The main difference is that unlike Peano and Hilbert they are not a single continuous line, and they are random rather than recursive. The Hausdorff dimensions of Peano and Hilbert are stated to be 2 [4], and by analogy one might expect that the fractal dimension is also 2 for the structures described here. This view, however, would lead to the conclusion that fractal D does not depend on the exponent c.

It is possible to make a crude estimate of the average distance d_{avg} from any point to the nearest point in any line. This is based on the idea[2] of area/length. We define

$$d_{avg} \approx \frac{1}{2} \frac{(boundary\ area)}{(total\ length\ of\ the\ lines)} \qquad (3)$$

From Eq. (2) we can estimate the total line length by approximating the sum of L_i by the corresponding integral. The total line length after placement n is proportional to $(N+n)^{-c+1}$ for large n and thus d_{avg} will be proportional to $(N+n)^{-1+c}$. A more precise determination of the average distance from any point to a point in the nearest line would probably yield the same functional form, namely that this distance tails off as $(N+n)^{-1+c}$. This steady decrease encourages us to think that these structures are space-filling in the limit.

It is hypothetically possible for all of the segments to lie in (say) the left-hand half of the bounding region, in which case the maximum distance from some points in the right-hand side to the nearest line segment will not diminish as i increases. While possible, such an arrangement is quite improbable (like all of the air molecules in a room going to the northern half of it). Perhaps the question of maximum distance from any point to the nearest line segment is not a true/false question, but one which can only be answered in terms of probabilities.

7. Formal Treatment of the Shrinking Maximum Length[3].

Let d_{max} be the greatest distance from *any* point not contained in any line to the nearest point in some line. We specifically consider the periodic boundary case of Fig. 2, so that boundaries do not enter the picture.

[2] One way to think of this is to imagine all of the line segments joined end-to-end and wound in serpentine fashion across the bounding area with a uniform pitch. This pitch will be approximately (boundary area)/(total line length).

[3] This analysis is due to Barry Cipra.

(Imagine the plane tiled with Fig. 2.) It is possible to prove that d_{max} tends to zero with probability 1. Imagine laying down a fine grid over the bounding region. At some point, the line segments will be short enough that any grid cell that doesn't already intersect with a line segment can completely contain the next segment, so eventually the random placement process is going to pick up that grid cell. (More precisely, the probability of not landing in it tends to zero.)

8. References.

[1] J. Shier and P. Bourke, Computer Graphics Forum, **32**, Issue 8, 89 (2013). This is the first publication in a refereed journal. The last version sent to the editor can be downloaded at [3].

[2] Benoit Mandelbrot, "Fractals: Form, Chance, and Dimension", W. H. Freeman, San Francisco (1977).

[3] J. Shier, web site http://john-art.com/stat_geom_linkpage.html

[4] A listing of fractals by Hausdorff dimension can be found in Wikipedia.

Appendix A -- Flow Chart.

The nonintersecting random search algorithm is most easily understood from a flow chart. It is really rather simple. The chart is readily adapted to units more complicated than a single line segment.