Available online at www.sciencedirect.com
SCIENCE DIRECT
Computers & Graphics 29 (2005) 998–1003
COMPUTERS & GRAPHICS
www.elsevier.com/locate/cag

Chaos and Graphics

Cantor knots

Cameron Browne

SWiSHzone.com Pty Ltd, The Basement, 33 Ewell St, Balmain 2041, Australia

Abstract

A technique based on Cantor division is used to embellish artistic knotwork designs. Weaving rules are relaxed to allow a mostly alternating weave to be applied to the resulting patterns. The new method is compared to traditional knotwork, and the generation of examples is discussed.

Keywords: Cantor set; Celtic knotwork; Alternating weave; Fractal; Art

1. Introduction

Celtic knotwork is an ornamental art style that has been practised for thousands of years. The knotwork designs found in illuminated religious manuscripts such as the Book of Kells from the 6th century AD and the Lindisfarne Gospels from the 8th century AD are widely considered to be its high point. The most important study of this art form is George Bain's book *Celtic Art: The Methods of Construction* [1].

Fig. 1 (left) shows a simple knot in the Celtic style. Typically, a single interlaced cord is used to trace a complex design to fill an area. Celtic knotwork is generally alternating in nature; tracing a path along a cord, the crossings will alternate over, under, etc. Shortly after the release of Bain's book, Thurston [2] proved that it is possible to impart an alternating weave to any set of interwoven curves that:

1. are finite and closed,
2. contain no points at which more than two cords cross, and
3. do not touch at any point without crossing.

These conditions hold for traditional Celtic knotwork designs.

Bain's passion for Celtic art was taken up by his son, Iain, who examined knotwork construction from a more geometrical perspective [3]. I. Bain demonstrates the technique of adding complexity to a design by splitting cords into parallel subcords that follow similar paths but with opposed parity at each crossing point *t*. For instance, Fig. 1 (left) shows a single-interlaced design and Fig. 2 (right) shows its double-interlaced counterpart. Note that wherever there exists an overpass on one subcord, then the matching crossing the same distance along the other subcord is an underpass, and vice versa. An alternating weave can always be imparted to a correctly *n*-interlaced design.

2. The Cantor set

The Cantor set is produced by taking the interval [0, 1], removing the middle third, removing the middle third of each of the two pieces thus created, and so on ad infinitum [4]. Fig. 2 (left) shows the Cantor set to five levels of recursion.

At level n of recursion, there are 2^n line segments, each of length $(2/3)^n$. The Capacity Dimension D (also called

E-mail address: Cameron.Browne@swishzone.com.

0097-8493/$ - see front matter
doi:10.1016/j.cag.2005.09.025

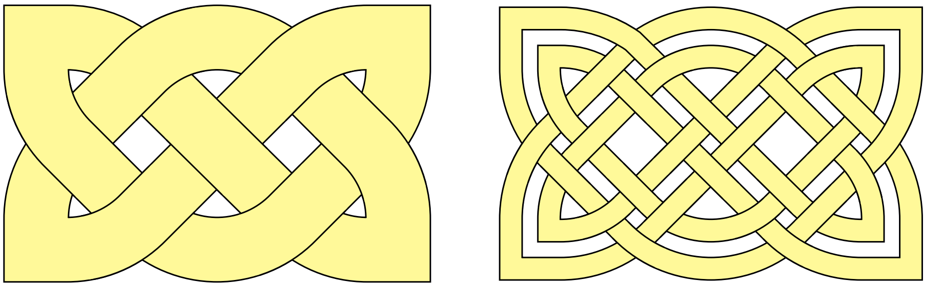

Fig. 1. A knot in the Celtic style and its double-interlaced counterpart.

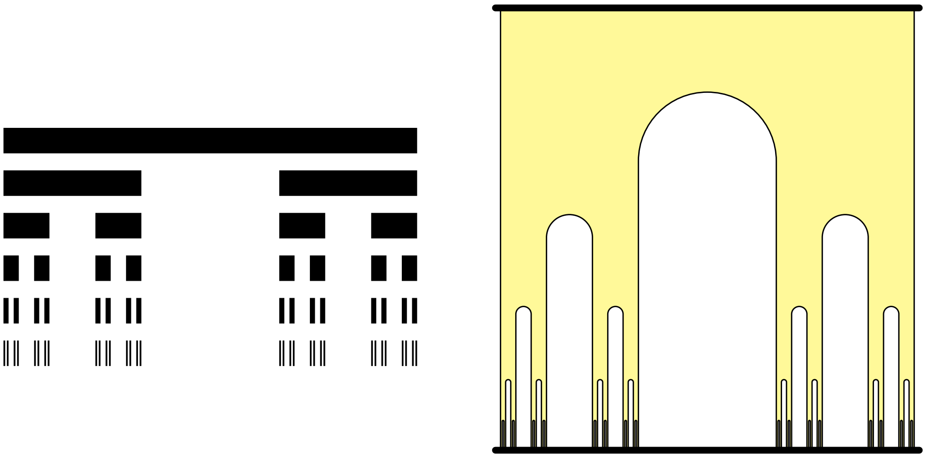

Fig. 2. The Cantor set and a Cantor braid.

the Fractal or Hausdorff Dimension) of the completely expanded Cantor set is given by

$$D = \ln 2/\ln 3 = 0.630929\ldots$$

The function CantorInterval( ) described in Listing 1 prints to file a line segment, then recursively applies the same operation to the left and right thirds of this segment. The C++ program shown in Listing 2 uses this function to write a complete Adobe PostScript 3.0 file "cantor.ps" that displays the Cantor set to a given level of recursion. The recursion level may be specified as a program argument (default is 8).

Code Listings

Listing 1. Recursive calculation of Cantor intervals to specific depth.

```
// CantorInterval( ): Prints a given cantor interval
// to file and iterates to the next level
//
// Parameters:
//   fp:        pointer to output file
//   level:     current level of iteration
//   max_level: maximum level of iteration (re-
//   cursion limit)
//   a, d:      end points of the current interval
//
void CantorInterval(FILE* fp, int level, int
max_level, float a, float d)
{
  fprintf(fp, "%f %f moveto %f %f lineto\n", a,
level + 0.5, d, level + 0.5);
  if (level < max_level)
  {
    float b = a + (d–a)/3.0;
    float c = a + (d–a)/1.5;
    CantorInterval(fp, level + 1, max_level, a, b);
    CantorInterval(fp, level + 1, max_level, c, d);
  }
}
```

Listing 2. Program to display the Cantor set to a specified depth in Adobe PostScript 3.0 format.

```
#include <stdio.h>
```

```
#include 〈stdlib.h〉
void main(int argc, char* argv[ ])
{
  int max_level = (argc<2) ? 8 : atoi(argv[1]);
  FILE* fp = fopen("cantor.ps", "w");
  if (fp != NULL)
  {
    fprintf(fp, "%%!PS-Adobe-3.0\n10 10 translate
\n500 50 scale\n\n");
    fprintf(fp, "0 setgray\n0.5 setlinewidth\n\nnew-
path\n");
    CantorInterval(fp, 0, max_level, 0, 1);
    fprintf(fp, "stroke\n\nshowpage\n");
    fclose(fp);
  }
}
```

3. Cantor braids and Cantor loops

The idea behind the Cantor set provides a novel method for introducing partial n-interlacement into a knotwork design. Fig. 2 (right) shows a single cord fixed at the top and split into successively smaller cords of 1/3 width that are fixed to the bottom. This construction may be called the Cantor braid, although it is not strictly a braid in the knot theory sense as there is a one-to-many mapping from top to bottom (see Sossinsky [5] for a good introduction to the field of knot theory).

Fig. 3 shows how the Cantor braid may be applied to a closed loop. An incision with diameter equal to 1/3 of the loop width is made along the loop's centerline. This incision may be stretched along the centerline; if completely stretched all the way around the loop then it will join itself to completely cut the loop into two parallel loops, each 1/3 the width of the original. This double loop is equivalent to that used by I. Bain for double-interlacement.

Fig. 4 shows some more complex incision patterns. Fig. 4 (left) shows a long level 1 incision accompanied by left and right level 2 incisions, while Fig. 4 (right) shows a number of periodic level 1 and level 2 incisions, demonstrating how multiple incisions may exist per level. For any interval I_n along which a level n incision occurs, there must also exist a parent incision at level $n-1$. In other words, it is not allowed to skip levels when making successively smaller incisions.

4. Weaving Cantor knots

Fig. 5 (left) shows a trivial knot that is isomorphic with the unknot. Fig. 5 (middle) shows the cord self-weaving through a single level 1 incision; this shape is no longer isomorphic with the unknot (note that each incision increases the shape's genus by 1). Fig. 5 (right) shows a more complex shape achieved by self-weaving

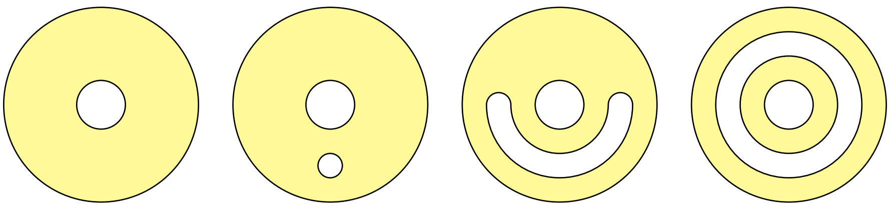

Fig. 3. A loop with a 1/3 incision stretched around to cut it into two loops.

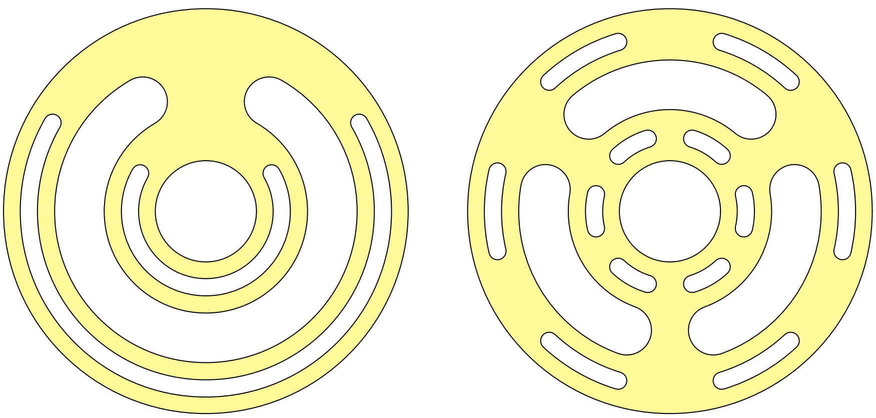

Fig. 4. More complex incision patterns.

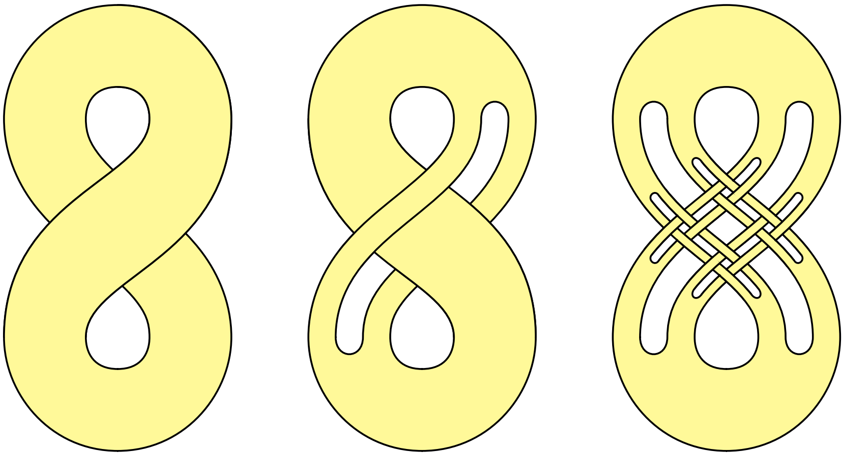

Fig. 5. A trivial knot, a Cantor knot with a single-woven level 1 incision, and a Cantor knot with multiply-woven level 2 incisions.

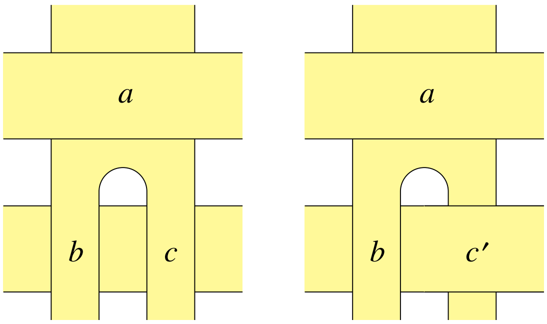

Fig. 6. Split weave rule: a bad weave (left) and an acceptable weave (right).

two level 1 and four level 2 incisions. Such designs with self-woven incisions shall be called Cantor knots.

Fig. 6 demonstrates how a mostly alternating weave can be imparted to Cantor knots. Firstly, a split is defined as a cord interval between consecutive crossings at different levels of incision. Fig. 6 (left) shows a bad weave that may result across a split. Following the vertical cord from a to b and from a to c, the alternating weave is maintained; however, alternation is lost when following the horizontal cord from b to c.

It is not possible to achieve perfect alternation both along and across a split at once. However, a reasonable solution can be achieved by alternating the weave across the split (i.e. in the direction of greatest weave frequency) and not caring about the weave along the split direction. This provides a degree of freedom that ensures that any Cantor knot can be woven successfully.

Fig. 6 (right) shows the split weave rule in action. The parity of crossing c has been inverted to c' so that the lowermost horizontal cord across the split now alternates from b to c'. The vertical cord alternates from a to b but no longer alternates along the split from a to c'. Since this non-alternation occurs along a size reduction, its visual impact is reduced and the resulting weave is acceptable. Crossings of opposite parity (b' and c) could have been used instead with a similar result.

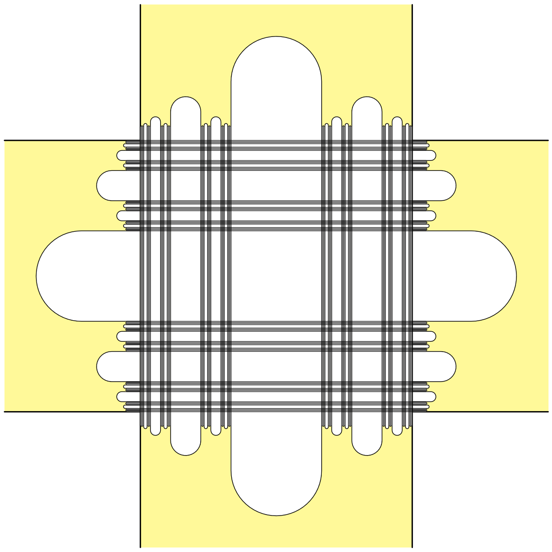

Fig. 7. Higher levels of recursion do not necessarily provide a more interesting weave.

One limitation of Cantor knots is that they only allow n-interlacement where n is a power of $2 = \{1, 2, 4, 8, 16, 32\ldots\}$, hence triple-interlacement is not possible. However, this is not a serious drawback, as triple

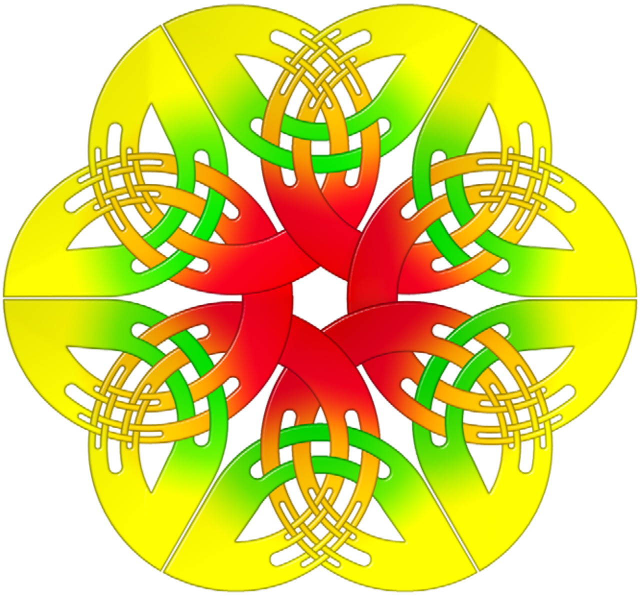
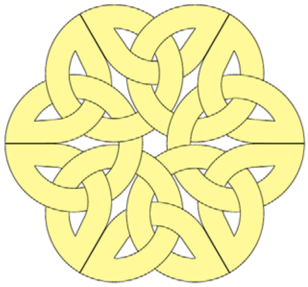

Fig. 8. An embellished knot woven at levels 0, 1 and 2, with source knot (inset).

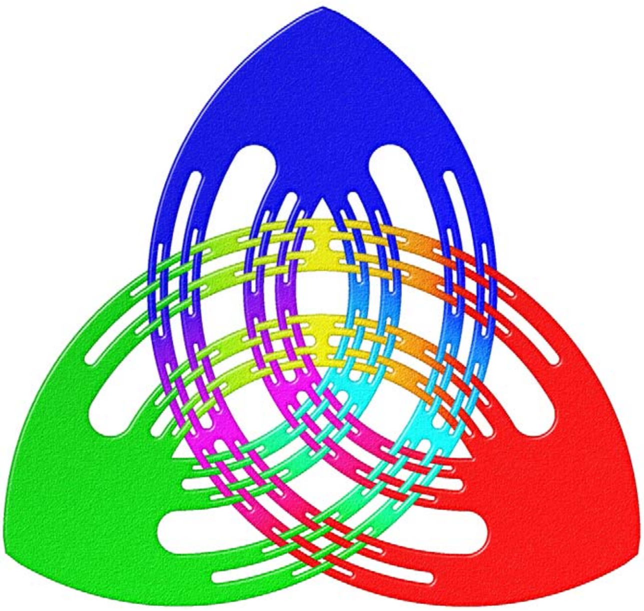
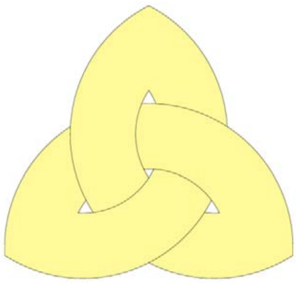

Fig. 9. An embellished knot woven at level 3, with source knot (inset).

interlacement and finer is rare in traditional Celtic knotwork. In fact, higher levels of division do not necessarily provide a more visually interesting weave, as demonstrated by Fig. 7; a design tends to lose some of its appeal once the weave becomes difficult to interpret.

5. Examples

Fig. 8 shows a decorative pattern based on these principles. Level 1 and level 2 incisions are applied to a source knot (inset) and self-woven. Three levels of weave occur: level 0 (i.e. full-width cord) near the center of the design, level 1 half-way out, and level 2 towards the outer periphery. The source knot is based on a traditional design found on a cross-slab stone in Ulbster, Caithness [1], adapted from its original square grid to a hexagonal one.

Fig. 9 shows a design with a simpler source shape (inset) but a more complex level 3 weave. This figure reinforces the point that a higher level of recursion does not necessarily produce a better pattern, and demonstrates that the finer the weave, the more basic the source shape must be. Note also that the higher-level weave leads to a flatter-looking result. The source knot, a trefoil, is a common feature in traditional knotwork, where it is typically used to fill small gaps.

The figures for this paper were created in Adobe PostScript 3.0 format. C++ code based on that shown in Listings 1 and 2 was used to generate filled shapes for most figures. The rendered full color images (Figs. 8 and 9) were then imported into Adobe Photoshop, and color, texture and shading filters applied.

6. Conclusion

This paper demonstrates how incisions made to knotwork designs can be self-woven to add complexity. The author knows of no traditional Celtic knotwork designs that employ this technique; it is a novel addition to the art form rather than a reconstruction of a known method.

Future work might include a more rigorous analysis of the split weave rule, and a formal proof that a given Cantor knot can always be woven successfully.

It may be interesting to reinterpret some traditional Celtic knotwork designs as Cantor knots of varying incision level and style. Of course, incisions do not have to lie exactly along the current cord's centerline, nor do they have to be exactly 1/3 of the current cord width; these are just natural and convenient default values that reinforce the conceptual relationship between Cantor knots and the Cantor set. It may also be interesting to weave parallel subcords along the cord path in a plait rather than across it.

References

[1] Bain G. Celtic art: the methods of construction. Glasgow: William McClellan and Co.; 1951.
[2] Thurston H. Celtic interlacing patterns and topology. Science News 1954;33:50–62.
[3] Bain I. Celtic knotwork. London: Constable; 1986.
[4] Weisstein E. Concise encyclopedia of mathematics. Boca Raton: CRC Press; 1999.
[5] Sossinsky A. Knots: mathematics with a twist. Cambridge: Harvard University Press; 2002.