Theory Comput. Systems **39**, 413–423 (2006)
DOI: 10.1007/s00224-005-1243-1

Theory of Computing Systems

More Fun with Symmetric Venn Diagrams*

Frank Ruskey and Mark Weston

Department of Computer Science, University of Victoria,
PO Box 3055, Victoria, British Columbia, Canada V8W 3P6
{ruskey,mweston}@cs.uvic.ca

Abstract. Many researchers have had fun searching for and rendering symmetric Venn diagrams, culminating in the recent result of Griggs, Killian, and Savage ("Venn Diagrams and Symmetric Chain Decompositions in the Boolean Lattice," *Electronic Journal of Combinatorics*, **11** (2004), R2), that symmetric Venn diagrams on n curves exist if and only if n is prime. Our purpose here is to point out ways to prolong the fun by introducing and finding the basic properties of Venn and near-Venn diagrams that satisfy relaxed notions of symmetry, while leaving tantalizing open problems. A symmetric Venn diagram is one that possesses a rotational symmetry in the plane. A monochrome symmetric Venn diagram is one that is rotationally symmetric when the colours of the curves are ignored. A necessary condition for the existence of a monochrome symmetric Venn diagram is that the number of curves be a prime power. We specify conditions under which all curves must be non-congruent and give examples of small visually striking monochrome symmetric Venn diagrams found by algorithmic searches. A Venn diagram partitions the plane into 2^n open regions. For non-prime n we also consider symmetric diagrams where the number of regions is as close to 2^n as possible, both larger and smaller.

Introduction

Venn diagrams and their close relatives, the Euler diagrams, form an important class of combinatorial objects which are used in set theory, logic, and many applied areas. Symmetric Venn diagrams are beautiful objects that have recently been intensively studied; however, only one notion of symmetry has been used. This naturally motivates the question: Can we expand on the traditional notions of symmetry in Venn diagrams?

* Frank Ruskey's research was supported in part by NSERC.

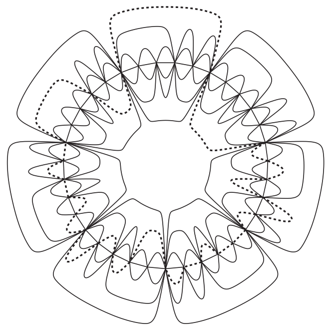

Fig. 1. Symmetric 7-Venn diagram with one curve highlighted (dashed).

Let $\mathcal{C} = \{C_1, C_2, \ldots, C_n\}$ be a family of n simple closed curves in the plane. The curves are required to be finitely intersecting. Let $X_i,\ i = 1, 2, \ldots, n$, be either the open bounded interior or the open unbounded exterior of the curve C_i. Then we say that $\mathcal{C}$ is a *Venn diagram* if all of the 2^n open regions $X_1 \cap X_2 \cap \cdots \cap X_n$ are non-empty and connected.[1] If the connectedness condition is dropped the diagram is called an *independent family*.

An i-region in a Venn diagram is a connected region that is interior to exactly i of the curves. A Venn diagram is *simple* if at most two curves intersect at any point.

An n-Venn diagram is *symmetric* if it consists of n congruent curves $C_1, C_2, \ldots, C_n$ where each successive curve is obtained from its predecessor by a rotation of $360i/n$ degrees about a fixed point x. We think of distinct curves as having distinct colours. Henderson [6] proved that a necessary condition for the existence of symmetric n-Venn diagrams is that n be prime. This condition is also sufficient: Griggs et al. recently gave a construction for symmetric Venn diagrams for any prime n [3]—a construction based on symmetric chain decompositions of the boolean lattice.

Given a ray emanating from x, a symmetric diagram may be partitioned into n sectors (sometimes referred to as "pie-slices") by taking the regions sliced by rotating the ray by $360i/n$ degrees for $i = 1, 2, \ldots, n$. The only difference between successive sectors is a permutation π of the curve colours. As observed by Ruskey [7], π must be a circular permutation.

Figure 1 shows a symmetric 7-Venn diagram with several interesting properties. First, consider that there are two adjacent vertices of degree $2n$ on the inner (and outer) face that isolate between them a sector comprising $1/n$th the entire diagram (here $n = 7$), and that this sector's structure repeats n times in the entire diagram. We can label the curves $1, 2 \ldots, n$ by arbitrarily fixing curve 1 (say the dotted curve in Figure 1) and adding one (modulo n) as the curve is rotated clockwise around the centre of the circle. The entire diagram can then be recovered by rotating the sector n times and rotating the label of each curve for each successive copy by adding one (modulo n). This corresponds to rotating the curve labels by applying the circular permutation $(1\ 2 \cdots n)$.

[1] Note that the term "Venn diagram" is often used informally to refer to diagrams with fewer than 2^n regions, for example in using them to illustrate set inclusion and exclusion, but here we exclusively use the term as defined. When fewer than 2^n regions occur, they are properly called *Euler diagrams*.

Previous research on generating symmetric diagrams has focused on generating such a sector of the diagram that can be rotated n times to create the entire diagram; see, for example, [1], [5], and [3].

1. Monochrome Symmetry

An n-Venn diagram is said to be (n, p)-monochrome symmetric if it possesses a p-fold rotational symmetry when the labels (i.e., colours) of the curves are ignored. Since the individual curves can be uniquely recovered from a simple Venn diagram, the concept is of interest only for non-simple diagrams.

Theorem 1. *In any (n, k)-monochrome symmetric Venn diagram, k is prime and n is a power of k.*

Proof. Each $1/k$ sector of the diagram contains the same number of i-regions. Thus k must divide $\binom{n}{i}$ for all $0 < i < n$, which in turn is equivalent to the condition that k is prime and n is a power of k. See, for example, Theorem 13 of [8]. $\square$

Figure 2 shows an example of the smallest non-trivial monochrome symmetric Venn diagram in which $n = 4$ and $k = 2$.

2. Pseudo-Symmetry

We now study a restricted version of monochrome symmetry. Consider the (4, 2)-monochrome Venn diagram shown in Figure 2: we see that the diagram maps onto itself under a rotation of 180 degrees about the centre point. Another way of viewing it is that the diagram is composed of two copies of the "pie-slice" forming the top half of the diagram, with the curves labelled in the second copy from the first by a permutation of the colours in the first pie-slice; we can generalize this to consider a monochrome symmetric diagram in which every pie-slice is formed by relabelling under the same permutation. This example leads us to make the following definition.

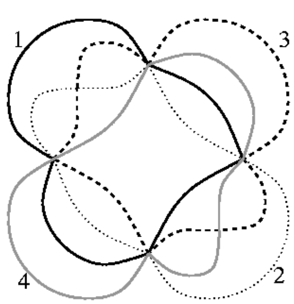

Fig. 2. A (4, 2)-monochrome Venn diagram.

Definition 1. An n-Venn diagram $\mathcal{V}$ has p-fold *pseudo-symmetry* with $1 < p < n$ if there is a pie-slice making up $1/p$th of $\mathcal{V}$ so that the rest of $\mathcal{V}$ can be generated by repeatedly rotating the pie-slice and recolouring the curves according to a fixed permutation π. We call π the *defining* permutation, and refer to a diagram with this type of symmetry as (n, p)*-pseudo-symmetric*.

The diagram in Figure 2 is (4, 2)-pseudo-symmetric, and the curves are labelled so that the permutation to be applied is (1234): since $p = 2$ the permutation is only applied once, to obtain the bottom half of the diagram from the top half.

This idea forms a natural extension of traditionally symmetric Venn diagrams: there, a single pie-slice forming $1/n$th of the diagram was rotated n times. Thus, we can generalize this idea to a Venn diagram with two parameters, namely n and p, and a similar process of generating the rest of the diagram from the pie-slice forming $1/p$th of the diagram; namely copying the pie-slice and rotating the curve labels. However, unlike symmetric Venn diagrams, pseudo-symmetric diagrams do not have the property that permuting the labels of the last pie-slice will yield the labelling of the first pie-slice since applying $(1\,2 \cdots n)$ exactly $p < n$ times will not return a curve label to itself.

This definition also implies that the diagram maps onto itself by a rotation of $360/p$ degrees ignoring the curve colourings, however, the extra condition about the permutation being fixed between successive copies of the pie-slice seems to be necessary in that there exist diagrams that map onto themselves by a rotation of $360/p$ degrees but do not use the same permutation between successive copies of the pie-slice before the last and the first: we refer to this condition as *weak pseudo-symmetry*. The diagram in Figure 3 has monochrome symmetry but not pseudo-symmetry, and thus is weakly pseudo-symmetric; Figure 4 illustrates that there are different permutations between successive pie-slices. Though an interesting idea, we do not consider weak pseudo-symmetry further.

The diagrams in Figures 2 and 3 also possess dihedral symmetry ignoring curve colourings, but this property in general will not apply to other pseudo-symmetric diagrams we construct.

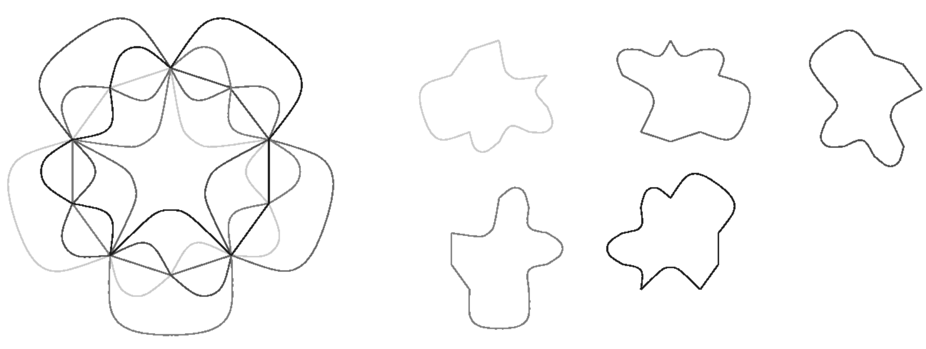

a) b)

Fig. 3. A monochrome symmetric but not pseudo-symmetric 5-Venn diagram with some curves congruent and some not; part (b) shows the separate curves.

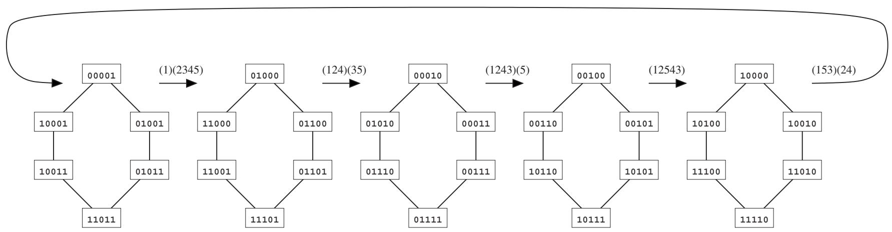

Fig. 4. The dual graph of the 5-Venn diagram in Figure 3 showing the distinct permutations between each pie-slice.

Note that we do not say that a normal symmetric n-Venn diagram is (n, n)-pseudo-symmetric, as this would lead to confusion when we consider the differences between pseudo-symmetry and normal symmetry.

3. Properties of Pseudo-Symmetric Diagrams

In considering symmetric Venn diagrams, one can use the fact that the permutation to be applied is always circular and can thus be assumed to be the permutation $(1\ 2 \cdots n)$ by relabelling the curves. For pseudo-symmetric diagrams, this is also true:

Lemma 1. *The defining permutation π of a pseudo-symmetric diagram must be a circular permutation.*

Proof. Consider labelling each region in the diagram by a bitstring x of length n, where x_i is 1 if the region is in the interior of curve i, and 0 otherwise. Since a diagram with (n, p)-pseudo-symmetry can be viewed as being composed of p copies of a pie-slice forming $1/p$th of the diagram, the $2^n - 2$ different labels of regions (ignoring the internal and external regions) must be divided into p sets of equal size, one per pie-slice. We can consider all sets but the first as being generated from the first by one or more applications of π to each region, viewed as a bitstring.

We assume π is not circular and derive a contradiction. If π is not circular, it can be written as $\pi = (\pi_1\pi_2 \cdots \pi_j)\pi'$, with $0 < j < n$ and π' is disjoint from $(\pi_1\pi_2 \cdots \pi_j)$. By relabelling we can assume that $\pi = (1\ 2 \cdots j)\pi'$. Any Venn diagram must contain a region corresponding to the bitstring $w = 1^j 0^{n-j}$, which is a fixed point of the action of π. However, each label in the non-final sets must have a corresponding label in the final set that they map onto via repeated applications of π. This implies that w is an extra element in that final set, as no other set can map onto it, contradicting the fact that the sets must have the same cardinality. $\square$

Since curve labellings are arbitrary, for convenience we can always relabel the curves so that the permutation π is $(1\,2 \cdots n)$. For example, the curves in Figure 2 are labelled so that the top half maps onto the bottom half by applying the permutation (1234).

The following lemma illustrates a major difference between pseudo-symmetry and symmetry in Venn diagrams.

Lemma 2. *In an (n, p)-pseudo-symmetric Venn diagram, if $p < n$, then no two curves are congruent.*

Proof. Assume the diagram is (n, p)-pseudo-symmetric, with $1 < p < n$. Consider the pie-slice forming $1/p$th of the entire diagram, which rotates p times to give the entire diagram, and the permutation $\pi = (1\,2 \cdots n)$ as given above that is applied to the curve labels each time to obtain the next pie-slice.

Consider the set of edges of a given colour, say colour 1, in this pie-slice. These edges form a connected path of that colour from one side to the other. We call this path a *curve segment*. There are n different curve segments passing through the pie-slice, one corresponding to each colour, and we can label these n segments by their colour in the first pie-slice: call s_i the segment with colour i. Each segment must be distinct in shape from all others. In the final diagram this pie-slice is rotated p times and the curve labels are incremented each time by applying the permutation π.

Thus, for example, in the second copy s_1 becomes colour 2 by applying π, in the third copy s_1 becomes colour 3, and in general in the jth rotation of the pie-slice, $1 < j \le p$, s_1 becomes colour $(j+1) \bmod n$. Moreover, segment s_i must be colour $(i+j-1) \bmod n$ in the jth copy, with $0 \equiv n$.

Consider the inverse permutation function π^{-1} defined by $\pi^{-1}[x] = y$ if and only if $\pi[y] = x$. In general the segment of colour i is s_i in the first copy, $s_{\pi^{-1}[i]}$ in the second copy (i.e., it is the segment whose colour shifts to colour i), so the curve segment of colour i in the jth copy of the pie-slice must be $s_{(i-j+1) \bmod n}$, with $1 \le i \le n$ and $1 \le j \le p$.

Now consider that the curve of colour i is made up of the cycle formed by joining all of the segments of colour i. Define the sequence of segments followed by curve i as the circular sequence $\mathcal{S}(i) = (s_i, s_{(i-1) \bmod n}, \ldots, s_{(i-p+1) \bmod n})$, with $1 \le i \le n$. Since $p < n$ it is clear that $\mathcal{S}(i) \ne \mathcal{S}(j)$ for $i \ne j$ given the permutation π. The sequence of segments of length p can be thought of as a "window" of size p moving around the circular list of segments, which has length n. Thus curves i and j are composed of a different circular sequence of distinct segments, and thus are distinct (not congruent). $\square$

As an example of Lemma 2, consider the $(4, 2)$-pseudo symmetric diagram in Figure 2. In the top half of the diagram that is rotated to obtain the bottom half, we can label the segments as black = 1, thin dotted = 2, thick dotted = 3, and light grey = 4. Then $\mathcal{S}(1) = (s_1, s_4)$ (i.e., the black curve follows segments 1 and then 4), $\mathcal{S}(2) = (s_2, s_1)$, $\mathcal{S}(3) = (s_3, s_2)$, and $\mathcal{S}(4) = (s_4, s_3)$. The four distinct curves are shown in Figure 5.

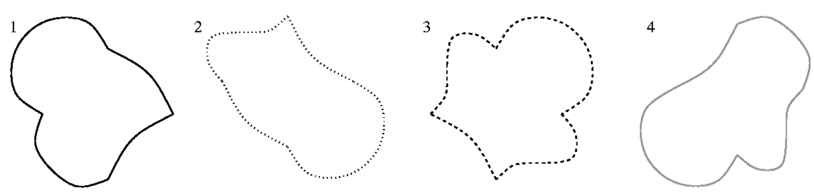

Fig. 5. The four distinct curves of the diagram in Figure 2.

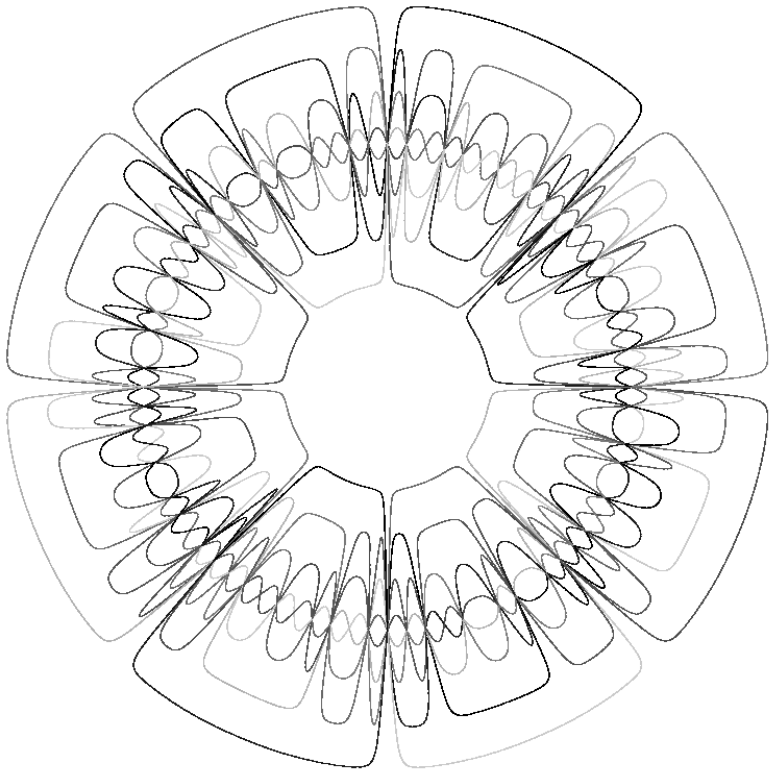

Fig. 6. An (8, 2)-pseudo-symmetric Venn diagram.

The Griggs, Killian, Savage (GKS) [3] result, mentioned earlier, involves generating symmetric Venn diagrams by creating the dual graph of the sector that is rotated *n* times to form the complete diagram. Both of the larger diagrams that follow were created by using an algorithm for the GKS construction and then modifying the output to ensure that it was a pseudo-symmetric Venn diagram.

An (8, 2)-pseudo-symmetric Venn diagram is shown in Figure 6. Though at first glance the diagram looks symmetric, upon closer inspection there are several differences between pie-slices that show that the diagram does in fact have only two-fold pseudo-symmetry.

Figure 7 shows a 9-Venn diagram with (9, 3)-pseudo-symmetry. Areas of the diagram that differ from sector to sector are magnified to clarify the differences. Further details of the construction of the two larger diagrams can be found in [9].

The smallest open case is $(16, 2)$, which is far too large to construct by hand: each sector forming $\frac{1}{16}$ of the diagram would contain 4096 regions.

An interesting open problem is to create a general construction for pseudo-symmetric diagrams, extending the GKS result.

Conjecture 1. *A* (p^k, p)*-pseudo symmetric Venn diagrams exist for all primes* p *and integers* $k > 1$.

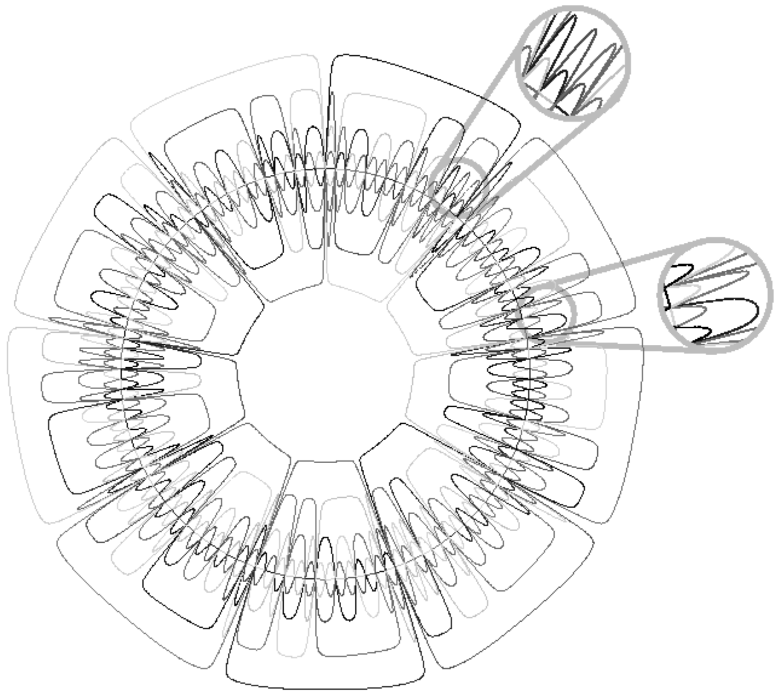

Fig. 7. A (9, 3)-pseudo-symmetric Venn diagram, with blown-up regions to illustrate the small differences between successive sectors.

4. Symmetric Non-Venn Diagrams

In the case that n is not prime, symmetric Venn diagrams cannot exist. Thus, it is natural to consider diagrams that are symmetric and in some sense "close" to Venn diagrams: if we can ensure the number of regions in a diagram is close to (either greater or lesser than) the 2^n required for an n-Venn diagram, we will in some sense "bracket" non-symmetric Venn diagrams of non-prime n both from above and below by symmetric (non-Venn) diagrams.

In [4] Grünbaum generalized the concept of symmetry in Venn diagrams to independent families of n sets. If n is not prime, there will be some i, $0 < i < n$, such that regions enclosed by i curves must be repeated to create a symmetric diagram, and thus we must have more than 2^n regions in any symmetric diagram. Grünbaum gave the lower bound $M(n)$ on the number of required regions: $M(n) = 2+n(C_n-2)$, where C_n is the number of binary necklaces of length n. The numbers C_n are well-studied, see pages 139–141 of [2] for explicit formulae and further references. We say that symmetric independent families that attain this lower bound are minimal, in the sense that they cannot have any regions deleted without either losing their symmetry or losing some regions altogether and ceasing to be independent families. Table 1 gives $M(n)$ for small values of n.

Grünbaum also conjectured in [4] that symmetric independent families with only $M(n)$ regions exist for every integer n (this includes the special case that n is prime, and

Table 1. $M(n)$ and $M'(n)$ for small values of n.

n	$M'(n)$	2^n	$M(n)$
4	14	16	18
6	56	64	74
8	242	256	274
9	506	512	524
10	1012	1024	1062
12	4070	4096	4202

the symmetric independent family is just a symmetric Venn diagram—it has no repeated regions), and gave examples for $n = 4$ and $n = 6$ of such diagrams. Figure 8 shows a symmetric independent family of eight curves with $M(8)$ regions. We have generated many such diagrams for non-prime n up to 10, see [9] for more examples.

In addition to the above approach of adding extra regions to create symmetric diagrams, we can also create them by removing as few regions as possible from a Venn diagram to create a diagram that is not an independent family (as some regions will be missing) but will retain rotational symmetry. Since these diagrams are not Venn

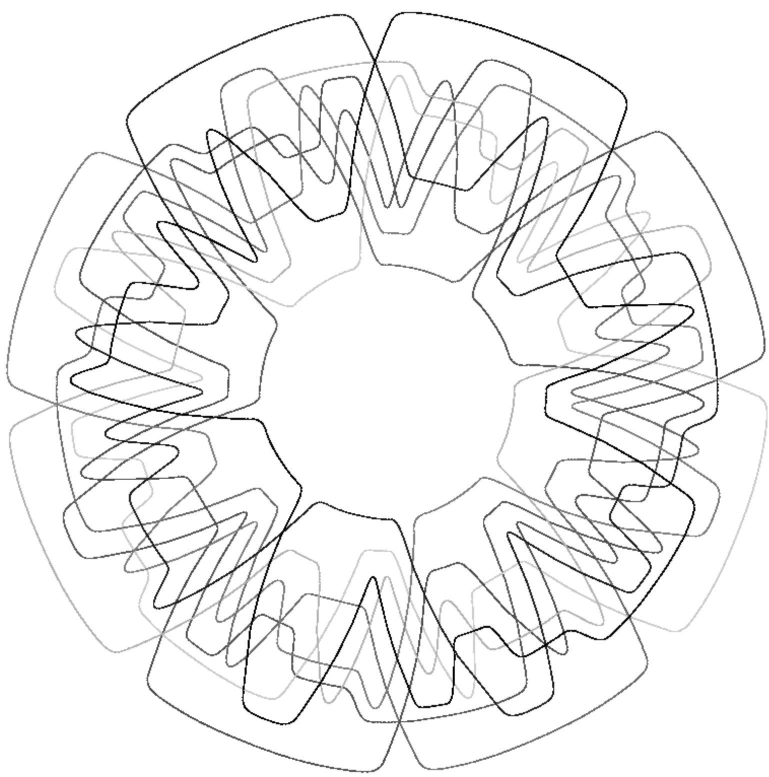

Fig. 8. A minimal symmetric independent family of eight curves.

diagrams but are "close" to being Venn diagrams in that they are missing few regions and every region is connected, we call them *near-Venn* diagrams. Symmetric near-Venn diagrams for non-prime n are maximally symmetric Euler diagrams in the sense that they cannot have any more regions added without either losing their symmetry or repeating regions.

Like the function $M(n)$, which provided a lower bound on the number of regions in a symmetric independent family of order n, we define a new function $M'(n)$, which provides an upper bound on the number of regions in a symmetric near-Venn diagram of order n with no repeated regions. An explicit expression is

$$M'(n) = 2 + n \sum_{k=1}^{n-1} \left\lfloor \binom{n}{k} / n \right\rfloor .$$

We give $M(n)$ and $M'(n)$ for small n in Table 1, and 2^n for comparison. For both symmetric independent families and for symmetric near-Venn diagrams the bound on the number of regions is asymptotic to 2^n.

An example of a near-Venn diagram generated in this fashion is shown in Figure 9. This diagram has $M'(6) = 56$ regions, eight fewer than the 64 required for a Venn diagram of order 6.

Figures 9 and 8 are also interesting in that they are as close as possible to being simple without being simple: the diagram in Figure 9 has six points where three or more curves intersect (every other point has the property, common to every point in simple diagrams, that at most two curves cross at that point). Grünbaum suggested in [4] that we cannot do better than this: he conjectured that if n is not prime, no simple symmetric independent family of n curves exists with $M(n)$ regions, and in fact we have been unable to find a simple example of either a symmetric independent family or near-Venn diagram for any n not prime. Thus, mirroring Grünbaum, we conjecture:

Conjecture 2. *A simple symmetric near-Venn diagram of n curves, with n non-prime, does not exist.*

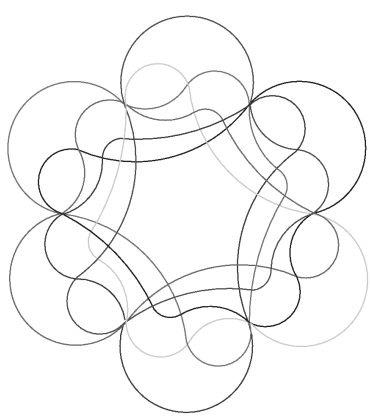

Fig. 9. A symmetric near-Venn diagram of six curves.

We have generated examples of maximal symmetric near-Venn diagrams for $n = 4, 6, 8, 9$ and more can easily be produced; in fact, most symmetric independent families can easily be turned into symmetric near-Venn diagrams by deleting regions. These diagrams were generated by a backtracking methodology; details on the generation process as well as pictures of the larger examples of both symmetric independent families and near-Venn diagrams can be found in [9]. An interesting area of future work is to find a general construction for symmetric independent families and near-Venn diagrams for n non-prime, and further explore their properties.

Acknowledgments

Many thanks to Branko Grünbaum for helpful discussion about the topics of this paper, and Stirling Chow for proofreading.

References

[1] A. Edwards. Seven-set Venn diagrams with rotational and polar symmetry. *Combinatorics*, *Probability*, *and Computing*, (7):149–152, 1998.

[2] R. Graham, D. Knuth, and O. Patashnik. *Concrete Mathematics*: *A Foundation for Computer Science*, second edition. Addison-Wesley, Reading, MA, 1994.

[3] J. Griggs, C. E. Killian, and C. Savage. Venn diagrams and symmetric chain decompositions in the boolean lattice. *The Electronic Journal of Combinatorics*, 11, Article R2, 30 pages (electronic), 2004.

[4] B. Grünbaum. The search for symmetric Venn diagrams. *Geombinatorics*, 8:104–109, 1999.

[5] P. Hamburger. Doodles and doilies, non-simple symmetric Venn diagrams. *Discrete Mathematics*, 257(2–3):423–439, 2002.

[6] D. Henderson. Venn diagrams for more than four classes. *American Mathematical Monthly*, 70(4):424–426, 1963.

[7] F. Ruskey. A survey of Venn diagrams. Dynamic survey. *The Electronic Journal of Combinatorics*, Article DS5, 2001.

[8] D. Singmaster. Divisibility of binomial and multinomial coefficients by primes and prime powers. In *A Collection of Manuscripts Related to the Fibonacci Sequence*, pages 98–113. Fibonacci, Santa Clara, CA, 1980.

[9] M. Weston. On symmetry in Venn diagrams and independent families. Master's thesis, University of Victoria, 2003.

Received July 16, 2004, *and in final form December* 1, 2004. *Online publication July* 13, 2005.

Note added in proof. In his Master's thesis, "Symmetric chain decompositions and independent families of curves" (Computer Science Dept., N.C. State University, 2003), Zongliang Jiang gives examples of symmetric independent families with $M(n)$ regions for $n \leq 16$, when n is composite.