

The Search for Simple Symmetric Venn Diagrams

Frank Ruskey, Carla D. Savage, and Stan Wagon

Many people are surprised to learn that Venn diagrams can be drawn to represent all of the intersections of more than three sets. This surprise is perfectly understandable since small Venn diagrams are often drawn with circles, and it is impossible to draw a Venn diagram *with circles* that will represent all the possible intersections of four (or more) sets. This is a simple consequence of the fact that circles can finitely intersect in at most two points and Euler's relation $F - E + V = 2$ for the number of faces, edges, and vertices in a plane graph. However, there is no reason to restrict the curves of a Venn diagram to be circles; in modern definitions a Venn diagram is a collection of simple closed Jordan curves. This collection must have the property that the curves intersect in only finitely many points and the property that the intersection of the interiors of any of the 2^n sub-collections of the curves is a nonempty connected region.

If a Venn diagram consists of n curves then we call it an n -Venn diagram. The *rank* of a region is the number of curves that contain it. In any n -Venn diagram there are exactly $\binom{n}{r}$ regions of rank r . Figure 1 shows a 2-Venn and two distinct 3-Venn diagrams. Note that the diagram in Figure 1(c) has three points where all three curves intersect. The regions in the diagrams of Figure 2 are colored according to rank.

The traditional three-circle Venn diagram has an appealing 3-fold rotational symmetry, and it is natural to ask whether there are n -Venn diagrams with an n -fold rotational symmetry for $n > 3$. Grünbaum [6] found a symmetric 5-Venn diagram made from ellipses. Henderson [10] noted the following necessary condition: if an n -Venn diagram has an n -fold

rotational symmetry, then n is prime. The reason is as follows:

In any symmetric n -Venn diagram the fixed point of the rotations, the center of the diagram, must lie in the unique region of rank n . The unbounded outer region has rank 0. Regions of rank $0 < r < n$ must be distributed symmetrically and thus their number, $\binom{n}{r}$, must be divisible by n . This property holds exactly when n is prime.

Why? Recall that $\binom{n}{r} = n(n-1) \cdots (n-r+1)/r!$. If n is prime and $0 < r < n$, then note that n occurs once in the right-hand side and all other numbers are less than n . On the other hand, if p is a nontrivial divisor of n , then the binomial coefficient with $r = p$ is the product of two integers $\binom{n}{p} = \frac{n}{p} \cdot m$ where $m = (n-1) \cdots (n-p+1)/(p-1)!$, but clearly p cannot divide m , and thus n does not divide $\binom{n}{p}$.

The elegant necessary condition of Henderson was long suspected to be sufficient, but it took some 40 years before it was proven to be sufficient by Griggs, Killian, and Savage [5]. In the intervening years, symmetric diagrams were discovered for n equal to 5, 7, and 11. Some of these diagrams are shown in Figure 2. The first symmetric 7-Venn diagrams were discovered independently by Grünbaum [7] and Edwards [3] (Fig. 2(b)); the first symmetric 11-Venn diagram was discovered by Hamburger [8].

A Venn diagram is said to be *simple* if exactly two curves pass through any point of intersection. The diagrams of Figures 1(a), (b) and 2(a), (b) are simple and the diagrams in Figures 1(c) and 2(c) are not simple. Simple Venn diagrams exist for all n , but no simple *symmetric* Venn diagrams are known for $n > 7$. On the other hand, no known condition precludes their existence for any prime n .

Venn diagrams were originally proposed as visual tools for representing "propositions and reasonings" [15] and how they are actually drawn in the plane will often influence how useful they are as tools. The definition of Venn diagram that we gave earlier is topological, but questions of geometry have also played a significant role in investigations of Venn diagrams. For example, one

Frank Ruskey is professor of computer science at the University of Victoria, Canada. His research is supported in part by NSERC. Carla D. Savage is professor of computer science at North Carolina State University. Her email address is savage@ncsu.edu. Her research is supported in part by NSF grant DMS-0300034. Stan Wagon is professor of mathematics at Macalester College. His email address is wagon@macalester.edu.




Figure 1: Symmetric n -Venn diagrams for $n = 2, 3$: (a) $n = 2$, (b) $n = 3$ simple, (c) $n = 3$ non-simple.




Figure 2: Symmetric Venn diagrams: (a) $n = 5$, (b) $n = 7$, (c) $n = 11$.

can ask: Which Venn diagrams can be drawn with all curves convex? For more than four sets, the practical usefulness of Venn diagrams diminishes but interesting mathematical questions arise. See [14] for a list of open problems related to Venn diagrams.

In this article we outline the technique of Griggs, Killian, and Savage [5] for producing symmetric Venn diagrams on a prime number of curves and the more recent efforts of Killian, Ruskey, Savage, and Weston [13] to create simple symmetric Venn diagrams. One of the diagrams from [13] was selected by Stan Wagon as the basis for the illustration shown on the cover; the method used to produce the image is described in the “About the Cover” description on page 1312.

Graph Theoretic Model

We first appeal to graph theory to get a “combinatorial” condition for Venn diagrams.

A Venn diagram D can be viewed as a (multi-)graph V embedded in the plane: the vertices of V are the points where curves of D intersect and the edges of V are the sections of the curves connecting the vertices. We can take the (geometric) dual of an

embedding of a planar graph V by placing a vertex v_r in every region r of V . If edge e separates regions r and s in V , then join v_r and v_s by an edge in the dual. The dual V^* of a Venn diagram is a planar embedding of a graph whose vertices are the subsets of $[n] = \{1, 2, \dots, n\}$.

To construct a Venn diagram, then, one could start with a graph whose vertices are the subsets of $[n]$. The n -cube Q_n is the graph whose vertices are the n -bit strings with edges joining strings that differ only in one bit. Since a subset $S \subseteq [n]$ can be viewed as an n -bit string whose i th bit is ‘1’ if and only if $i \in S$, the vertices of Q_n are in one-to-one correspondence with the regions in a Venn diagram. But Q_n is not planar for $n \geq 4$, so we cannot produce a Venn diagram simply by taking the dual of Q_n .

There is a theorem in graph theory that says: *In a planar graph G , if S is a bond, that is, a minimal set of edges whose removal disconnects G , then the edges in the dual G^* , corresponding to those in S , form a cycle in G^* .* For a proof, see West [16, Theorem 6.1.14]. This is exactly what is needed. If G^* is to be a Venn diagram, then for each $1 \leq i \leq n$, the graph G^* must have a corresponding cycle C_i




Figure 3: (a) A symmetric chain decomposition in \mathcal{B}_4 ; (b) embedding with cover edges, with each edge colored by its type.




Figure 4: (a) An overlay with the dual of the graph in Figure 3(b); (b) the resulting Venn diagram for 4 sets (the two vertices with arrows are identified).

to separate the sets containing i from those that do not. The dual of C_i back in G will be the set of edges joining vertices representing sets that do contain i to those that do not, and *this must be a bond of G .*

A spanning subgraph of Q_n is called *monotone* if every n -bit string with k ones is adjacent to a string with $k - 1$ ones (if $k > 0$) and to a string with $k + 1$ ones (if $k < n$). In a monotone subgraph of Q_n , for each $1 \leq i \leq n$, the edges joining vertices with i th bit ‘1’ to those with i th bit ‘0’ form a bond. Thus the following condition on a spanning subgraph G of Q_n will guarantee that the dual of G is a Venn diagram: *G is planar and monotone.* It is worth noting that this condition is not necessary; there are Venn diagrams for which G is not monotone.

In the next section, we show how to build a planar, monotone, spanning subgraph of Q_n using a *symmetric chain decomposition* in the Boolean lattice.

The Combinatorics

Return to the Boolean lattice \mathcal{B}_n whose elements are the subsets of $[n]$, ordered by inclusion. The Hasse diagram of \mathcal{B}_n , regarded as a graph, is isomorphic to

Q_n . A *chain* in \mathcal{B}_n is a sequence $S_1 \subseteq S_2 \subseteq \dots \subseteq S_t$ of elements of \mathcal{B}_n such that $|S_i| = |S_{i-1}| + 1$. The chain is *symmetric* if $|S_1| + |S_t| = n$. A *symmetric chain decomposition* of \mathcal{B}_n is a partition of the elements of \mathcal{B}_n into symmetric chains.

A significant result in order theory is that \mathcal{B}_n has a symmetric chain decomposition for every $n \geq 0$. One construction, due to Greene and Kleitman [4], works as follows. Regard the elements of \mathcal{B}_n as n -bit strings. View ‘1’ bits as right parentheses and ‘0’ bits as left parentheses and in each string, match parentheses in the usual way. This process may leave some ‘1’ or ‘0’ bits unmatched. From every string x with no unmatched ‘1’ grow a chain as follows. Change the first unmatched ‘0’ in x to ‘1’ to get its successor, y . Change the first unmatched ‘0’ in y (if any) to ‘1’ to get its successor. Continue until a string with no unmatched ‘0’ is reached. The chains shown in Figure 3(a), built using this rule, give a symmetric chain decomposition of \mathcal{B}_4 .

These chains form a planar spanning subgraph of Q_n . But to make the subgraph monotone, we need to add edges (without destroying planarity) to “cover” the first and last elements of each chain.

The chain starting at x can be covered by the chain starting at y where y is obtained from x by changing the last '1' in x to '0'. Not only do x and y differ in one bit, but so do the last elements of these chains. Viewing y 's chain as the parent of x 's chain, it can be shown that a preorder layout of the tree of chains gives a planar embedding of the chains together with their cover edges. A planar embedding of the subgraph of Q_4 consisting of the chains and the cover edges is shown in Figure 3(b).

The dual graph is shown in Figure 4(a). Say that an edge in the graph of Figure 3(b) has *type* i if it joins vertices that differ in position i . In Figure 4(a), a dual edge is colored according to the type of the edge it crosses. Figure 4(b) shows the resulting Venn diagram.

This method gives yet another constructive proof that for every $n \geq 0$, Venn diagrams exist for n sets. (A similar construction is implicit in [2], although they make no mention of symmetric chains.) So what about rotational symmetry? As described earlier, this is not possible if n is composite. But when n is prime, we can extract ideas from the construction described here to achieve symmetry, as shown in the next section.

Rotational Symmetry When n is Prime

When n is prime, the idea for constructing a rotationally symmetric Venn diagram is to somehow work within "1/ n -th" of \mathcal{B}_n (or Q_n) to get "1/ n -th" of the Venn diagram embedded in a "1/ n -th" pie-slice of the plane and then rotate the result by $2\pi i/n$ for $1 \leq i < n$ to complete the diagram. Fortunately, when n is prime, \mathcal{B}_n comes with a natural partition into symmetric classes.

For $x = x_1x_2 \cdots x_n$, let σ denote the left rotation of x defined by $\sigma(x) = x_2x_3 \cdots x_nx_1$. Let $\sigma^1 = \sigma$, and $\sigma^i(x) = \sigma(\sigma^{i-1}(x))$, where $i > 1$. Define the relation Δ on the elements of \mathcal{B}_n by $x \Delta y$ if and only if $y = \sigma^i(x)$ for some $i \geq 0$. Then Δ is an equivalence relation that partitions the elements of \mathcal{B}_n into equivalence classes called *necklaces*. When n is prime, every n -bit string, other than 0^n and 1^n , has n distinct rotations, so its necklace has exactly n elements.

In the hope of adapting the method of the previous section, we ask: *When n is prime, is there a way to choose a set R_n of necklace representatives, one from each necklace, so that the subposet of \mathcal{B}_n induced by R_n , $\mathcal{B}_n[R_n]$ has a symmetric chain decomposition?* Fortunately, the answer is *yes* (see next section), so construction of a rotationally symmetric Venn diagram can proceed as follows.

Start with the strategically chosen set R_n of necklace representatives. Let $Q_n[R_n]$ be the subgraph of Q_n induced by R_n . The symmetric chain decomposition in $\mathcal{B}_n[R_n]$, together with appropriate cover edges, gives a planar, spanning, monotone subgraph P of $Q_n[R_n]$. Embed P in a 1/ n -th pie

slice of the plane with 1^n at the center and 0^n at the point at infinity. Note that, as graphs, $Q_n[R_n]$ and $Q_n[\sigma^i(R_n)]$ are isomorphic for any bitwise rotation σ^i of the vertices. So $Q_n[\sigma^i(R_n)]$ has a subgraph P_i isomorphic to P . Then rotating the embedding of P by $2\pi i/n$ about the origin gives a planar embedding of P_i . Taken together, the embeddings of the P_i give a rotationally symmetric planar embedding of a spanning monotone subgraph of Q_n and therefore the dual is a Venn diagram. Finally, the dual is drawn in a symmetric way. The entire process is illustrated for $n = 5$ in the sequence of Figures 5(a), (b), (c), (d). The chains in $Q_5[R_5]$ are 10000-11000-11100-11110 and 10100-10110 (see the lowest "hexagon" in Fig. 5(a)).

Choosing Necklace Representatives

Here is a way to choose a set R_n of necklace representatives, one from each necklace, so that the subposet of \mathcal{B}_n induced by R_n has a symmetric chain decomposition.

Define the *block code* $\beta(x)$ of a binary string x as follows. If x starts with '0' or ends with '1', then $\beta(x) = (\infty)$. Otherwise, x can be written in the form:

$$x = 1^{a_1}0^{b_1}1^{a_2}0^{b_2} \cdots 1^{a_t}0^{b_t}$$

for some $t > 0$, where $a_i > 0$, $b_i > 0$, $1 \leq i \leq t$, in which case,

$$\beta(x) = (a_1 + b_1, a_2 + b_2, \dots, a_t + b_t).$$

As an example, the block codes of the string 1110101100010 and all of its rotations are shown below.

bit string	block code	bit string	block code
1110101100010	(4, 2, 5, 2)	1100010111010	(5, 2, 4, 2)
0111010110001	(∞)	0110001011101	(∞)
1011101011000	(2, 4, 2, 5)	1011000101110	(2, 5, 2, 4)
0101110101100	(∞)	0101100010111	(∞)
0010111010110	(∞)	1010110001011	(∞)
0001011101011	(∞)	1101011000101	(∞)
1000101110101	(∞)		

When n is prime, no two different rotations of an n -bit string can have the same *finite* block code. Assuming that block codes are ordered lexicographically, in each necklace of n -bit strings (except 0^n , 1^n) the unique string with minimum block code can be chosen as the representative.

For n prime, let R_n be the set of n -bit strings that are the minimum-block-code representatives of their necklaces. Build chains using the Greene-Kleitman rule, except chains start with a string with exactly one unmatched '1' and end at a string with exactly one unmatched '0'. Note that a node x and its successor y have the *same* block code, so if x has the minimum block code among all of its rotations, then so does y . Thus every element of x 's chain is the (minimum-block-code) representative of its necklace.

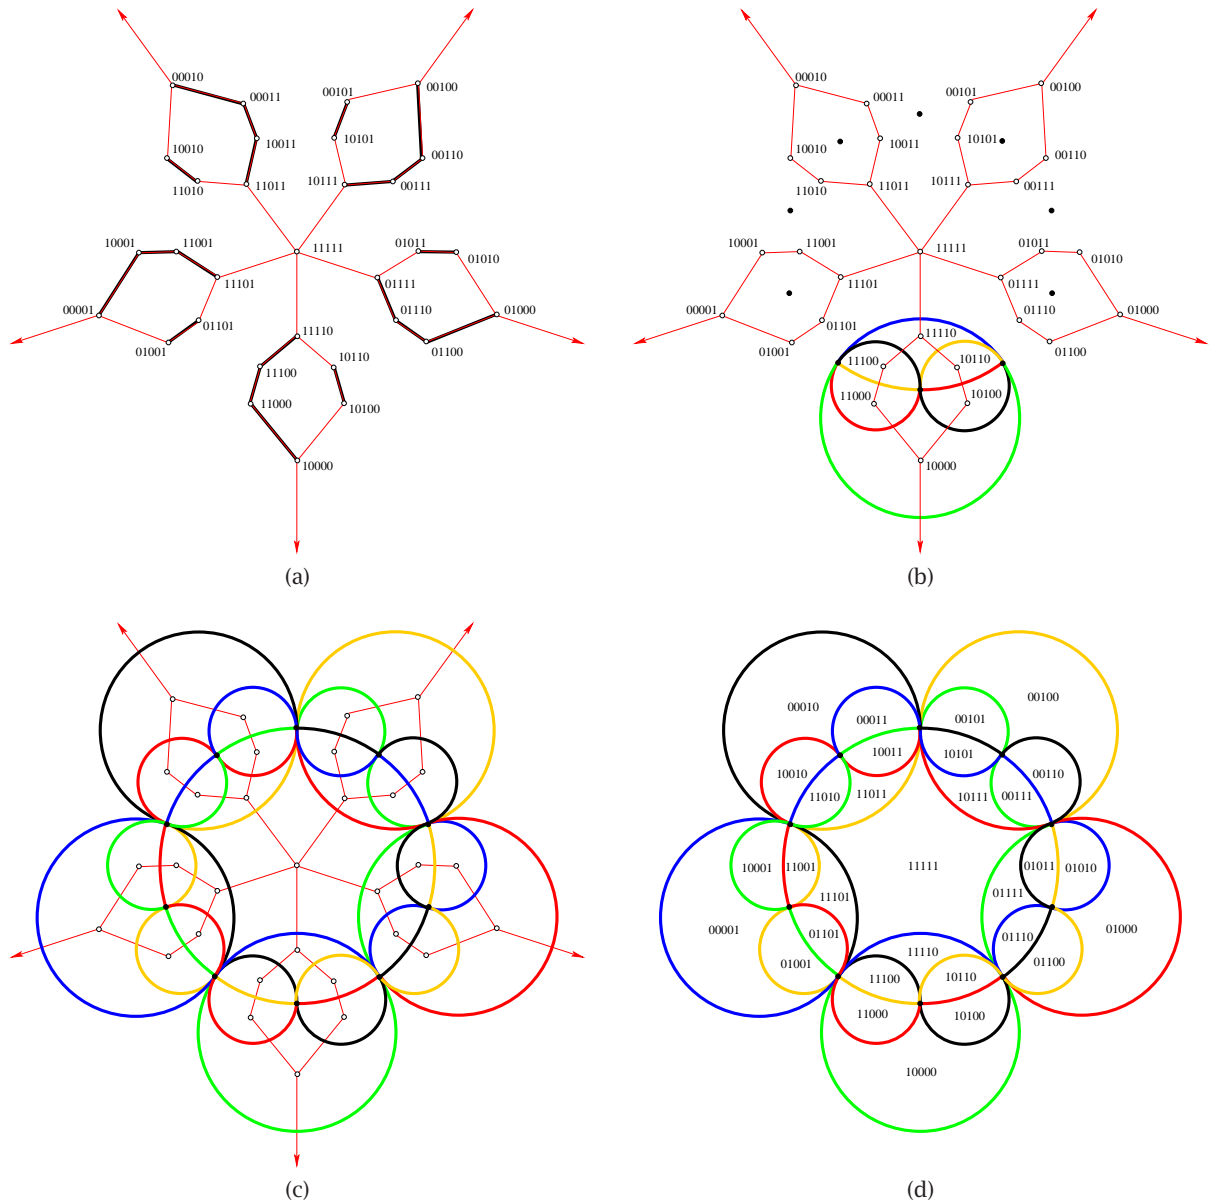


Figure 5: Constructing a symmetric 5-Venn diagram: (a) dual with symmetric chains highlighted, (b) the curves corresponding to the first chain cover, (c) repeating in each sector, (d) the final Venn diagram.

Simpler Venn Diagrams and Euler's Formula

Recall that a Venn diagram is *simple* if at most two curves intersect at any given point. This means that, viewed as a graph, every vertex of a simple Venn diagram has degree 4. The number of faces is 2^n , since every subset of $[n]$ corresponds to a region, and the number of edges is half the sum of the vertex degrees, so by $V - E + F = 2$, a simple Venn diagram has $2^n - 2$ vertices. In contrast, the number of vertices in the Venn diagrams we have constructed via symmetric chain decompositions is the number of

elements in the middle levels of \mathcal{B}_n : $\binom{n}{\lfloor n/2 \rfloor}$, which is roughly $2^n / \sqrt{n}$. This means that the average number of curves intersecting at any given point is about $c\sqrt{n}$ for some constant c . But a hidden feature of the Greene–Kleitman symmetric chain decomposition will allow a dramatic improvement.

Since the number of faces of a Venn diagram is fixed and since $V - E + F = 2$, once $E > V$, more vertices in the diagram mean the average degree is smaller and thus, on average, fewer curves intersect at any point. If the Venn diagram is the dual of a planar spanning monotone subgraph G of Q_n

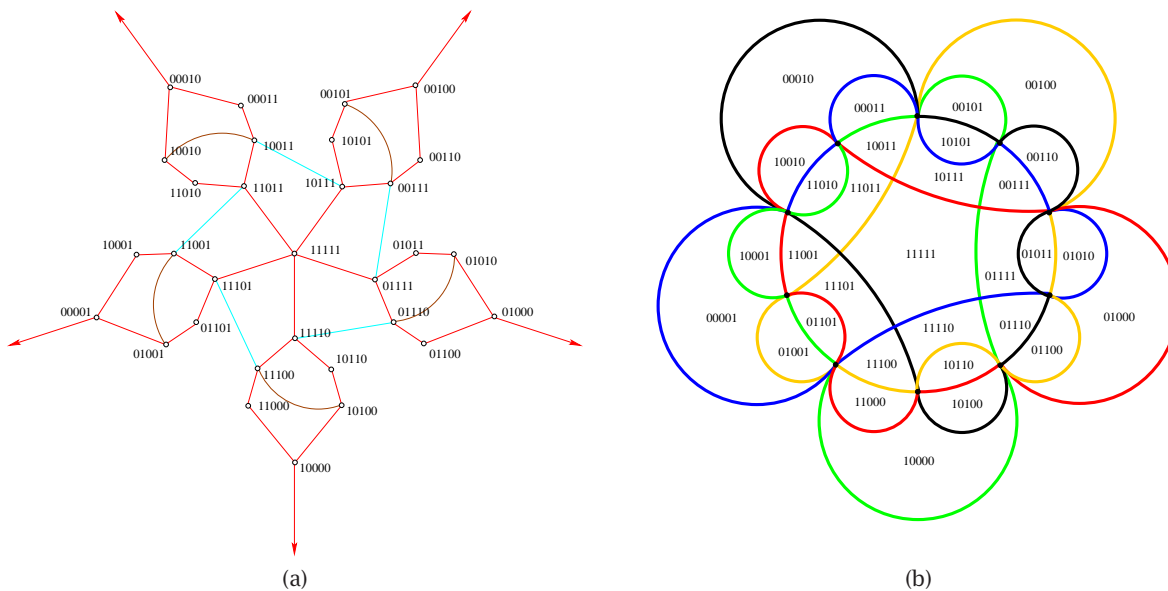


Figure 6: (a) The dual with two simplifying edges added in pie slice. (b) The effect of the cyan simplifying edge (compare with Fig. 5 (d)) is to increase the number of vertices from 10 to 15.

that has been embedded in the plane, we can increase the number of vertices of the Venn diagram by increasing the number of faces of G . One way to do this is to add edges of Q_n to G , without destroying the planarity of G . The added edges join vertices which differ in one bit. For example, Figure 6 shows the addition of ten simplifying edges to the 5-Venn dual of Figure 5 and the effect that adding the five cyan ones has on the resulting 5-Venn diagram. Note that the number of vertices increases from 10 to 15.

The Greene-Kleitman symmetric chain decomposition provides a systematic way to do this: *Any face bounded by two chains and two (suitably chosen) cover edges can always be “quadrangulated” by edges joining vertices that differ in one bit.* This is illustrated in Figures 7 and 8. Furthermore, it can be shown that as $n \rightarrow \infty$, the number of vertices in the resulting Venn diagram is at least $(2^n - 2)/2$, so the average number of curves intersecting at any given point is at most 3. Since $(2^n - 2)/2$ is half the number of vertices in a simple Venn diagram, the diagrams of [13] were dubbed “half-simple”. (Experiments suggest that as $n \rightarrow \infty$, the construction is doing better than 50%, perhaps closer to 60%.)

The construction is certainly not optimal. Figure 9 shows that further simplifying edges of Q_n , beyond those specified by the construction, can be added. To date the simplest symmetric 11-Venn diagram is due to Hamburger, Petruska, and Sali [9]; their diagram has 1837 vertices and is about 90% simple.

Figure 9 was the starting point for the half-simple Venn diagram shown on the cover. Figure

10(a) shows the result of embedding the graph of Figure 9 in a “1/11th” pie slice of the plane and then rotating it by $2\pi i/11$ for $1 \leq i < 11$ to get a monotone, planar, symmetric, spanning subgraph of Q_{11} . Its dual, drawn by Wagon’s *Mathematica* program and shown in Figure 10(b), is a half-simple, symmetric 11-Venn diagram. The program regards (a) as a planar map, so the regions have been 4-colored to highlight this interpretation.

Figure 11(a) shows one curve of the 11-Venn diagram. Each of the 11 curves is a rotation of this one. Figure 11(b) shows the Venn diagram with regions colored by rank and with one curve highlighted.

Acknowledgements: We are grateful to Mark Weston and Nikolay Dinev for helpful discussions and to Mark for creating Figure 2(c). Thanks also to the referees for helpful comments. Support from the National Sciences and Engineering Research Council of Canada and the National Science Foundation is acknowledged and appreciated.

References

- [1] BARRY CIPRA, PETER HAMBURGER, and EDIT HEPP, Aesthetic aspects of Venn diagrams, *Proceedings of the 2005 Bridges Conference on Mathematical Connections in Art, Music and Science, July 31–August 3, 2005*, Banff, Canada, 339–42.
- [2] BETTE BULTENA and FRANK RUSKEY, Venn diagrams with few vertices, *Electronic Journal of Combinatorics*, 5, R44, (1998), 21 pages.
- [3] ANTHONY W. F. EDWARDS, Seven-set Venn diagrams with rotational and polar symmetry, *Combinatorics, Probability, and Computing*, 7 (1998), 149–52.




Figure 7: Symmetric chain decomposition (a), with cover edges (b), and quadrangulated faces (c), in $Q_7[R_7]$.




Figure 8: Symmetric chain decomposition with cover edges and quadrangulated faces in $Q_{11}[R_{11}]$ from the construction of [13].






Figure 9: After manually adding extra edges, including wrapping edges, to the graph of Figure 8.

- [4] CURTIS GREENE and DANIEL J. KLEITMAN, Strong versions of Sperner's theorem, *Journal of Combinatorial Theory, Series A*, **20** (1976) 80–8.
- [5] JERROLD GRIGGS, CHARLES E. KILLIAN, and CARLA D. SAVAGE, Venn diagrams and symmetric chain decompositions in the Boolean lattice, *Electronic Journal of Combinatorics*, **11** (1), R2, (2004), 30 pages.
- [6] BRANKO GRÜNBAUM, Venn diagrams and independent families of sets, *Mathematics Magazine*, **48** (Jan-Feb 1975) 12–23.
- [7] ———, Venn diagrams II, *Geombinatorics*, **2** (1992), 25–32.
- [8] PETER HAMBURGER, Doodles and doilies, non-simple symmetric Venn diagrams, *Discrete Mathematics*, **257** (2002), 423–39.
- [9] P. HAMBURGER, GY. PETRUSKA, and A. SALI, Saturated chain partitions in ranked partially ordered sets, and non-monotone symmetric 11-Venn diagrams, *Studia Scientiarum Mathematicarum Hungarica*, **41** (2004), 147–91.
- [10] D. W. HENDERSON, Venn diagrams for more than four classes, *American Mathematical Monthly*, **70** (1963), 424–26.
- [11] JOAN P. HUTCHINSON, Three coloring Siamese trees, Personal communication, 2006.
- [12] JOAN P. HUTCHINSON and STAN WAGON, Kempe revisited, *American Mathematical Monthly*, **105** (1998), 170–74.




(a)




(b)

Figure 10: (a) The plane graph, derived from Figure 9, whose dual is a half-simple 11-Venn diagram (with regions 4-colored). (b) All 11 curves of the half-simple 11-Venn diagram created by taking the dual of the graph in (a).



(a)



(b)

Figure 11: (a) One curve of the half-simple 11-Venn diagram. (b) The half-simple 11-Venn diagram, with regions colored according to rank, and one curve highlighted.

[13] CHARLES E. KILLIAN, FRANK RUSKEY, CARLA D. SAVAGE, and MARK WESTON, Half-simple symmetric Venn diagrams, *Electronic Journal of Combinatorics*, 11 (1), R86, (2004), 22 pages.

[14] FRANK RUSKEY and MARK WESTON, A survey of Venn diagrams, *Electronic Journal of Combinatorics*, 4 DS5, (1997), (updated 2001, 2005).

[15] JOHN VENN, On the diagrammatic and mechanical representation of propositions and reasonings, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 9 (1880), 1-18.

[16] DOUGLAS B. WEST, *Introduction to Graph Theory*, Second edition, 2001.