



Generating simple convex Venn diagrams

Khalegh Mamakani, Wendy Myrvold, Frank Ruskey*

Department of Computer Science, University of Victoria, BC, Canada

ARTICLE INFO

Article history:

Available online 30 April 2012

Keywords:

Spherical Venn diagram

Symmetry

Exhaustive enumeration

Grünbaum encoding

ABSTRACT

In this paper we are concerned with producing exhaustive lists of simple monotone Venn diagrams that have some symmetry (non-trivial isometry) when drawn on the sphere. A diagram is simple if at most two curves intersect at any point, and it is monotone if it has some embedding on the plane in which all curves are convex. We show that there are 23 such 7-Venn diagrams with a 7-fold rotational symmetry about the polar axis, and that 6 of these have an additional 2-fold rotational symmetry about an equatorial axis. In the case of simple monotone 6-Venn diagrams, we show that there are 39020 non-isomorphic planar diagrams in total, and that 375 of them have a 2-fold symmetry by rotation about an equatorial axis, and amongst these we determine all those that have a richer isometry group on the sphere. Additionally, 270 of the 6-Venn diagrams also have the 2-fold symmetry induced by reflection about the center of the sphere.

Since such exhaustive searches are prone to error, we have implemented the search in a couple of ways, and with independent programs. These distinct algorithms are described. We also prove that the Grünbaum encoding can be used to efficiently identify any monotone Venn diagram.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

Named after John Venn (1834–1923), who used diagrams of overlapping circles to represent propositions, Venn diagrams are commonly used in set theory to visualize the relationships between different sets. The familiar three circle Venn diagram is usually drawn with a three-fold rotational symmetry (Fig. 1(a)) and the question naturally arises as to whether there are other Venn diagrams with rotational and other symmetries. Some research has been done recently on generating and drawing Venn diagrams of more than three sets, particularly in regard to symmetric Venn diagrams, which are those where rotating the diagram by $360/n$ degrees results in the same diagram up to a relabeling of the curves. Grünbaum [10] discovered a rotationally symmetric 5-Venn diagram (Fig. 1(b)). Henderson [12] proved that if an n -curve Venn diagram has an n -fold rotational symmetry then n must be prime. Recently, Wagon and Webb [17] clarified some details of Henderson's argument. The necessary condition that n be prime was shown to be sufficient by Griggs, Killian and Savage [9] and an overview of these results was given by Ruskey, Savage, and Wagon [14].

A Venn diagram is *simple* if at most two curves intersect at any point. In this paper we only consider simple Venn diagrams. There is one simple symmetric 3-Venn diagram and one simple symmetric 5-Venn diagram. Edwards wrote a program to exhaustively search for simple polar symmetric 7-Venn diagrams and he discovered 5 of them, but somehow overlooked a 6-th [7]. His search was in fact restricted to monotone Venn diagrams, which are equivalent to those that can be drawn with convex curves [2]. Fig. 1(c) is a 7-set Venn diagram with 7-fold rotational symmetry, called “Adelaide” by Edwards, and which was discovered independently by Grünbaum [11] and Edwards [7]. It should be noted that the diagrams

* Corresponding author.

E-mail addresses: aahmadi@cs.uvic.ca (K. Mamakani), wendy@csc.uvic.ca (W. Myrvold), ruskey@csc.uvic.ca (F. Ruskey).

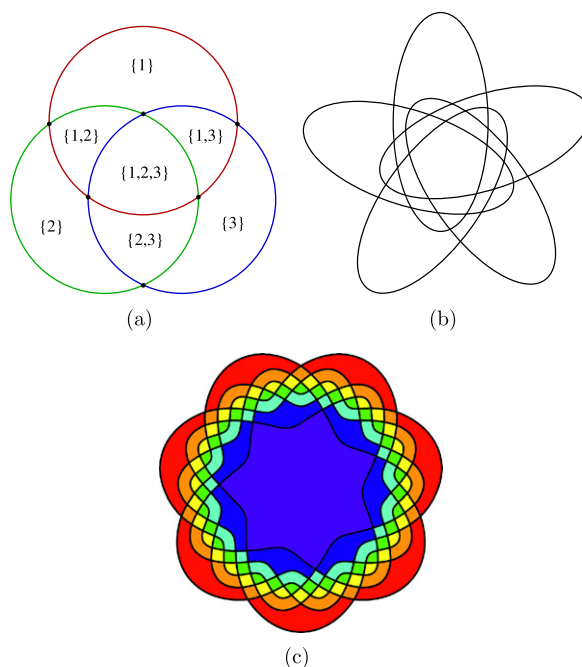


Fig. 1. (a) A 3-Venn diagram whose curves are circles. (b) A 5-Venn diagram whose curves are ellipses. (c) A symmetric 7-Venn known as “Adelaide.”

constructed in [9] are inherently non-simple, and the existence question for simple symmetric 11-Venn diagrams remains an open problem.

It is known that Venn diagrams exist for any number of curves and several constructions of them are known [15], but the total number of simple Venn diagrams is known only up to $n = 5$. In this paper, we determine that the number of simple monotone 6-Venn diagrams is 39 020; undoubtedly there are many other non-monotone diagrams.

Symmetric spherical Venn diagrams were first systematically investigated by Weston [18] and the recent paper [16] shows that Venn diagrams exhibiting each of the possible order 2 isometries exist for all n . The underlying constructions of [16] are inherently non-simple and the diagrams presented in the current paper are the first known simple Venn diagrams with certain order 2 isometries for 6-Venn diagrams.

A program was written to search for monotone simple symmetric 7-Venn diagrams and 23 of them were reported in the original version of the “Survey of Venn Diagrams” (Ruskey and Weston [15]) from 1997, but no description of the method was ever published and the isomorphism check was unjustified. Later Cao [3] checked those numbers, and provided a proof of the isomorphism check, but again no journal paper with the result was ever published. In this paper, we justify that isomorphism check with a new simpler proof and yet again recompute and verify the number of symmetric simple 7-Venn diagrams.

In this paper we are restricting our attention to the special (and most studied) class of Venn diagrams; diagrams that are both simple and monotone (drawable with convex curves). Our eventual aim is to provide a complete enumeration of such diagrams for small values of n , determining also the diagrams which have non-trivial isometries when embedded on the sphere. The underlying techniques rely on exhaustive backtrack searches with intelligent pruning rules whose use is justified by structural theorems. Such computer searches are prone to error and so we have made considerable effort to ensure that our computations are correct by using different representations, different methods for checking isomorphism, and independent programming efforts.

Our main concern in this paper are Venn diagrams with n curves, where $n = 6$ and $n = 7$. The case of $n = 7$ is done first, but only on those diagrams that have an order 7 rotational symmetry, because of the overwhelming number of possibilities otherwise. Two different representations are used, with three independent programs. We find that there are 23 non-isomorphic rotationally symmetric simple monotone 7-Venn diagrams. Of these 23, there are 6 that have an additional 2-fold “polar symmetry” (Fig. 10), and 17 that do not (Fig. 11).

In the case of $n = 6$ we again used 2 different representations and three independent programs. There are 39 020 non-isomorphic simple monotone 6-Venn diagrams. Of these, 375 have polar symmetry, of those 27 have an isometry group order of 4, and 6 have an isometry group order of 8 (Fig. 12). Additionally, 270 of the 6-Venn diagrams also have the 2-fold symmetry induced by reflection about the center of the sphere.

We introduce several different representations of these diagrams. Although these representations are somewhat similar in nature and there are efficient algorithms for getting from one representation to the other, we used them to implement independent generating algorithms for each class of studied Venn diagrams.

The remainder of this paper is organized as follows. In Section 2 we introduce the terminology and basic definitions. In Section 3 we explain various representations of Venn diagrams. The generating algorithms and results are described in Section 4.

Some of the results of this paper were first published in the conference proceedings [13] and [4]. In this paper we have combined, extended, and refined those results.

2. Definitions

Let $\mathcal{C} = \{C_0, C_1, \dots, C_{n-1}\}$ be a collection of n finitely intersecting simple closed curves in the plane. We call \mathcal{C} an *independent family* if each of 2^n sets

$$X_0 \cap X_1 \cap \dots \cap X_{n-1}$$

is nonempty, where X_i is either the unbounded open exterior or open bounded interior of curve C_i . If each set is a nonempty and connected region, then \mathcal{C} is called an *n-Venn diagram*. A *simple Venn diagram* is one in which exactly two curves cross each other at each intersection point.

A *k-region* in a Venn diagram is a region that is in the interior of precisely k curves. In an *n-Venn diagram*, each *k-region* corresponds to a *k-element* subset of a set with n elements. Thus, there are $\binom{n}{k}$ distinct *k-regions*. A Venn diagram is *monotone* if every *k-region* is adjacent to both some $(k-1)$ -region (if $k > 0$) and also to some $(k+1)$ -region (if $k < n$). The *rank* of a region of a Venn diagram is defined to be $\sum_{i=0}^{n-1} 2^i x_i$, where $x_i = 1$ if it is in the interior of curve i and $x_i = 0$ otherwise. By the definition of a Venn diagram, each region has a unique rank r in the range $0 \leq r < 2^n$.

An *n-Venn diagram* is *rotationally symmetric*, if rotation of the diagram by an angle of $2\pi/n$ about a fixed point in the plane does not change the diagram, except for a relabeling of the curves. Therefore, a $1/n$ th circular sector of a rotationally symmetric *n-Venn diagram* is enough to generate the whole diagram.

Polar symmetry is another type of symmetry, which was introduced by Grünbaum. Consider a Venn diagram as being projected onto a sphere with the rank 0 and rank $2^n - 1$ regions mapped to the north and south poles. A *polar flip* is the rotation of sphere by π radians about an equatorial axis; thus the northern and southern hemispheres are exchanged. The polar flip of a plane Venn diagram is then obtained by projecting the flipped sphere back onto the plane – this has the effect of interchanging the “insides” and “outsides” of all the curves. A Venn diagram is *polar symmetric* if it can be drawn so that it is invariant under some polar flip.

Given a planar Venn diagram V , let V_M be its mirror image, let V_P be its polar flip, and let V_{MP} be the mirror image of its polar flip. On the plane we define two Venn diagrams V and V' to be *isomorphic* if V' can be changed into V , V_M , V_P , or V_{MP} by a continuous transformation of the plane (and thus the combinatorial structure is preserved). Note that the natural definition of Venn diagram isomorphism usually does not include the polar flips and we broaden the definition here to allow for polar flips as well.

A Venn diagram is *convex* if it is isomorphic to a Venn diagram with all curves drawn convexly. It has been proven that a Venn diagram is convex if and only if it is monotone [2]. In this paper we use the term convex and monotone interchangeably. Fig. 2 shows a simple monotone 6-Venn diagram. If each curve of a Venn diagram touches the outermost region, then it is called an *exposed Venn diagram*. In a convex (monotone) Venn diagram, every 1-region is adjacent to the empty region. Therefore, every convex Venn diagram is exposed.

We will refer to a Venn diagram that is embedded on the sphere as a *spherical Venn diagram*. Spherical Venn diagrams can potentially exhibit a richer set of symmetries than planar Venn diagrams. There are three types of isometries of the sphere. First, there is *rotation* about an axis; the polar flip is an example of such an isometry. Secondly, there is *reflection* about a plane. Finally, there is *rotary reflection*, which is a rotation followed by a reflection across a plane orthogonal to the rotational axis. A particularly simple type of rotary reflection is obtained when the rotation is by π radians. In that case, each point is mapped to the corresponding point on the opposite side of the sphere; we refer to this isometry as *antipodal symmetry*.

It is often easier to visualize a spherical Venn diagram by using its *cylindrical representation* where the surface of the sphere maps to a rectangle in the plane. Consider a Venn diagram that has been projected onto the surface of a sphere with radius r . A *cylindrical projection* of the Venn diagram can be obtained by mapping the surface of the sphere to a $2\pi r$ by $2r$ rectangle on the plane, where the equator of the sphere maps to a horizontal line of length $2\pi r$ and the north and south pole of the sphere are mapped to the top and bottom sides of the rectangle respectively. In this representation, the top of cylinder is usually assumed to represent the empty region and the bottom of the cylinder represents the innermost region. Fig. 2 shows a polar symmetric 6-Venn diagram and its cylindrical representation and Fig. 3 is the cylindrical representation of a 6-Venn diagram with a rotary reflective symmetry (in particular, it has antipodal symmetry).

Depending on the face that we choose as the empty region, we can project a diagram on the sphere to different diagrams on the plane. We will say that a spherical Venn diagram is *monotone* if it has some monotone projection. However, other projections of a monotone diagram are not necessarily monotone; in fact, every spherical *n-Venn diagram* with $n > 2$ has some non-monotone planar projection. Note that, unlike the monotone Venn diagrams on the plane, it is not known whether every spherical monotone Venn diagram is convexly-drawable or not.




Fig. 2. A simple monotone 6-Venn diagram and its cylindrical representation. The diagram is polar symmetric and the black dots show where the equatorial axis intersects the surface of the sphere. A flip by π radians about that axis leaves the diagram fixed, up to a relabeling of the curves.




Fig. 3. Cylindrical representation of a 6-Venn diagram with antipodal symmetry (a type of rotary reflection).

2.1. Venn diagrams as graphs

A Venn diagram can be viewed as a plane graph where the intersection points of the Venn diagram are the vertices of the graph and the sections of the curves that connect the intersection points are the edges of the graph. Thinking of a Venn diagram as a graph has many benefits and will provide us with one of our fundamental representations. In this representation the faces of the graph are the regions of the diagram. We will use either term, depending on our point of view, but note the following. By an r -region we mean a region that is on the interior of r curves; by *size of a face* we mean the number of edges that bound the face and by a k -face we mean a face of size k .

For a plane graph with f faces, v vertices and e edges, Euler's formula states that $f + v = e + 2$. The graph of an n -Venn diagram has 2^n faces. In a simple Venn diagram each vertex of this graph has degree 4; i.e. $e = 2v$, so a simple n -Venn diagram has $2^n - 2$ vertices (i.e., intersection points). The graph of a Venn diagram has the following properties.

Lemma 1. A simple Venn diagram on three or more curves is a 3-connected graph [5].

Lemma 2. There are no two edges in a face of a Venn diagram that belong to the same curve [5].

Lemma 3. In a simple Venn diagram on three or more curves there are no faces of size two [5].

Lemma 4. In a simple Venn diagram with more than three curves, there are no two faces of size 3 adjacent to another face of size 3.

Proof. Suppose there is a Venn diagram \mathcal{V} that has two 3-faces adjacent to another 3-face. Then as we can see in Fig. 4, there are two faces (the shaded regions) in the diagram with the same rank which contradicts the fact that \mathcal{V} is a Venn diagram. \square

An embedding of graph $G(V, E)$ on a surface S is a mapping τ of G to S such that:




Fig. 4. A single 3-face adjacent to two other 3-faces, the shaded regions have the same rank.

- For each vertex $v \in V$, $\tau(v)$ is a distinct point of S , i.e. the mapping is injective;
- The edges of G are mapped to disjoint open arcs of S ;
- For any edge $e = (u, v)$, $\tau(e)$ joins the points $\tau(u)$ and $\tau(v)$;
- For any edge $e = (u, v)$ and any vertex x where $x \neq u$ and $x \neq v$, $\tau(e)$ does not include $\tau(x)$.

The complement of $\tau(G)$ relative to surface S is a set of *regions* or *faces* in S . The embedding is called a *2-cell embedding* if every face is homeomorphic to an open disk.

For a graph $G(V, E)$, each edge $e = (u, v) \in E$ has two oriented directions which are represented as two *half-edges* $(u, v)_e$ and $(v, u)_e$ that are referred to as *twins*. Let \mathcal{E} be the set of all half-edges of graph G . A rotation system of G is a pair (σ, ϕ) where both σ and ϕ are permutations of \mathcal{E} . For each $a \in \mathcal{E}$, $\sigma(a)$ is the next half-edge in clockwise order in the circular list of half-edges incident to the same vertex and $\phi(a)$ is the twin of a . We usually describe a rotation system by the circular lists of incident edges of all vertices. For a given graph G , it has been proven that each rotation system uniquely describes a 2-cell embedding of G on some orientable surface S [6]. A *planar graph* is a graph that can be embedded on the sphere (equivalently on the plane). A *plane graph* is an embedded planar graph. The *mirror* of a plane graph is a plane graph obtained by reversing the circular list of incident edges with each vertex.

We follow Brinkmann and McKay's [1] definitions regarding isomorphism of plane graphs. Let $G_1 = (V_1, E_1, C_1)$ and $G_2 = (V_2, E_2, C_2)$ be two plane graphs, where C_1 and C_2 indicate the rotation systems of G_1 and G_2 respectively. We say G_1 is isomorphic to G_2 , if there is a bijective mapping ϕ from (V_1, E_1) to (V_2, E_2) that preserves the combinatorial structure; that is, if (e_1, e_2, \dots, e_k) is the circular list of edges in E_1 incident to $v \in V_1$, then $(\phi(e_1), \phi(e_2), \dots, \phi(e_k))$ is the circular list of edges in E_2 incident to $\phi(v) \in V_2$.

3. Representing Venn diagrams

In this section we introduce the representations that are used when we generate simple monotone Venn diagrams. First we introduce the Grünbaum encoding and we prove that each Grünbaum encoding identifies a simple exposed Venn diagram up to isomorphism. In the second part we discuss the binary matrix representation where each 1 in the matrix represents an intersection point of the corresponding Venn diagram. Having the matrix representation of a diagram, it is easy to check if it is a Venn diagram or not. In the third part we show how to represent simple monotone Venn diagrams using integer compositions. We use this representation to generate all polar-symmetric convex 6-Venn diagrams. Finally, we discuss representing simple convex Venn diagrams using a finite sequence of exchanges of curve labels. We generate all simple convex 6-Venn diagrams using this method.

3.1. Grünbaum encoding

Grünbaum encodings were introduced by Branko Grünbaum as a way of hand-checking whether two symmetric Venn diagrams are distinct [15]. We generalize this concept here to all Venn diagrams, symmetric or not, and then focus on the special properties that they have when the diagram is symmetric. The Grünbaum encoding of a simple exposed Venn diagram consists of $4n$ strings, 4 for each curve C_i . Call the strings w_i, x_i, y_i, z_i for $i = 0, 1, \dots, n-1$. In fact, given any one of the w, x, y, z strings of a Venn diagram V , we can obtain the other three. However, we need these four strings to compute the lexicographically smallest string as the Grünbaum encoding representative of V . Given the lexicographically smallest Grünbaum encoding of two Venn diagrams, then we can check if they are isomorphic or not.

Starting from one of the curves in the outermost or innermost regions, we first label the curves from 0 to $n-1$ in the clockwise or counter-clockwise direction. The starting curves of these labelings are chosen arbitrarily, and thus there can be several Grünbaum encodings of a given Venn diagram. Table 1 indicates whether the labeling starts on the inside or outside and whether the curve is considered to be oriented clockwise or counter-clockwise. To get the w_i strings we arbitrarily pick a curve and label it 0. It intersects the outer face in exactly one segment; the remaining curves are labeled $1, 2, \dots, n-1$ in a clockwise direction. Now that each curve is labeled, we traverse them, recording the curves that each intersects, until it returns, back to the outer face. Thus w_i is a string over the alphabet $\{0, 1, \dots, n-1\} \setminus \{i\}$. The strings x_i, y_i, z_i are produced in a similar manner, except that we are starting on the inner face, or traversing in a counter-clockwise direction, or both, as indicated in Table 1(a). In Table 1(b) we show part of the Grünbaum encoding of Fig. 1(b).

Table 1

(a) Conventions for Grünbaum encoding. (b) Grünbaum encoding of the 5-Venn diagram of Fig. 1(b).

(a)			(b)	
	CW	CCW	w:	1, 3, 2, 4, 1, 4, 2, 4, 1, 3, 1, 4
Outermost	w	y	x:	1, 4, 2, 4, 1, 3, 1, 4, 1, 3, 2, 4
Innermost	x	z	y:	1, 4, 2, 4, 1, 3, 1, 4, 1, 3, 2, 4
			z:	1, 3, 2, 4, 1, 4, 2, 4, 1, 3, 1, 4

For curve i , each string of Grünbaum encoding starts with $(i + 1)$ and ends with $(i + n - 1) \bmod n$. Since each one of w, x, y or z uses a different labeling of the same curves, there are permutations that map the labelings of one to the labelings of other. Given a Grünbaum encoding $\{w, x, y, z\}$ let the permutations π, σ and τ map the curve labels of w to the curves labels of x, y, z respectively. Let ℓ_i denote the length of string w_i and let $w_i[k]$ be the k th element of w_i where $k = 0, 1, \dots, \ell_i - 1$. We can get $y_{\sigma(i)}$ by

$$y_{\sigma(i)}[k] = \sigma(w_i[\ell_i - k - 1]).$$

To obtain $x_{\pi(i)}$ and $z_{\tau(i)}$, we first determine the unique location p in w_i where all curves have been encountered an odd number of times (and thus we are now on the inner face). We then have

$$x_{\pi(i)} = \pi(w_i[(k + p) \bmod \ell_i]),$$

and

$$z_{\tau(i)} = \tau(w_i[(k - i) \bmod \ell_i]).$$

In the case of a rotationally symmetric Venn diagram we only need the four strings w_0, x_0, y_0, z_0 to specify the Grünbaum encoding, since the others will be a trivial relabeling of these. E.g., for any other curve $i \neq 0$, we can get w_i by $w_i = w_0 + i \bmod n$. The other three strings of curve i can be obtained in the same manner from the corresponding strings of curve 0.

The following lemma gives the length of each string of the Grünbaum encoding of a simple symmetric Venn diagram.

Lemma 5. Each string of the Grünbaum encoding of a simple symmetric n -Venn diagram has length $(2^{n+1} - 4)/n$.

Proof. Clearly each string will have the same length, call it L . Recall that a simple symmetric n -Venn diagram has $2^n - 2$ intersection points. By rotational symmetry every intersection point represented by a number in the encoding corresponds to $n - 1$ other intersection points. However, every intersection point is represented twice in this manner. Therefore, $nL = 2(2^n - 2)$, or $L = (2^{n+1} - 4)/n$. \square

Let g be a Grünbaum string, i.e. $g \in \{w, x, y, z\}$. Each intersection point of curves i and j is represented by an entry of value j in g_i and an entry of value i in g_j . So each element of g_i of value j uniquely corresponds to an element of g_j of value i and vice versa. We call the corresponding elements of g_i and g_j *twins*. Let $g_i[k] = j$ be an intersection of curve i with curve j . For any curve c other than i and j , let η_c be the number of occurrences of c in g_i up to $g_i[k]$, starting from the first element of g_i . For each of the four Grünbaum strings of curve i , the parity of η_c shows whether $g_i[k]$ is in the interior or exterior of curve c . For example, for $g = w$ or $g = y$, if η_c is odd then $g_i[k]$ is in the interior of curve c and if η_c is even then $g_i[k]$ is in the exterior of curve c . The *weight* of $g_i[k]$ is defined to be the number $\text{weight}(g_i[k]) = (r_{n-1} \cdots r_1 r_0)_2$ where $r_k = 0$ if $k \in \{i, j\}$ and otherwise $r_k = \eta_k \bmod 2$.

Lemma 6. Let g be a Grünbaum string of an n -Venn diagram, where $n \geq 3$. For each pair of curves (i, j) , if $g_i[k] = j$ for some k , then there is a unique index l such that $\text{weight}(g_j[l]) = \text{weight}(g_i[k])$.

Proof. Since $g_i[k] = j$, there is a corresponding intersection point P where i and j intersect. Thus, when following curve j we will also encounter P , and so there must be an l such that $g_j[l] = j$.

To show uniqueness, we argue by contradiction, and assume that for some $m \neq l$, there is another entry in g_j such that $\text{weight}(g_j[m]) = \text{weight}(g_j[l])$. This entry must correspond to a second point of intersection P' of curves i and j . Let R be the region of the Venn diagram that is interior to exactly the same set of curves as P and P' , and let r be the rank of R . Thus both P and P' must be on the boundary of R . Thus, by Lemma 2, R must be a 2-face. But this is a contradiction, since Lemma 3 states that there are no 2-faces in a Venn diagram if $n \geq 3$. \square

Unlike the corresponding theorem in [4], the theorem below holds for all simple Venn diagrams, monotone or not.

Theorem 7. Given a Grünbaum encoding G of a simple exposed n -Venn diagram V , we can recover V from G , up to isomorphism of Venn diagrams.

Proof. It is known that a plane embedding of a 3-connected planar graph is unique, once the outer face has been identified [19]. We will present a constructive proof which shows that the Grünbaum encoding determines a 2-cell embedding of the diagram on the sphere. Assume that the given Grünbaum encoding is $\{w_i, x_i, y_i, z_i\}_{i=1}^n$. We use a three step algorithm to construct the rotation system that uniquely represents the Venn diagram. In the first two steps we associate a vertex label with each $w_i[k]$ for all i and k , and then based on those labels, we create the rotation system in step 3.

- **Step one:** Starting with w_0 , for each $w_i[k]$ with $w_i[k] > i$, we associate a new vertex label with $w_i[k]$. At the end of this process there are $2^n - 2$ distinct vertex labels since every intersection occurs exactly twice in w . At the end of this step vertex labels have been assigned to all intersections of curves i and j , where $0 \leq i < j \leq n - 1$.
- **Step two:** We now associate vertex labels with the remaining entries of w ; but we must be careful to provide the correct label, since the same pair of curves can intersect multiple times. Let v be the vertex label associated with $j = w_i[k]$ where $i < j$. We need to uniquely locate the value of ℓ such that $i = w_j[\ell]$ is the twin of $w_i[k]$. By Lemma 6 there will be a unique value of ℓ such that $\text{weight}(w_i[k]) = \text{weight}(w_j[\ell])$, which can be determined by a simple scan of w_j . We then associate v with $w_j[\ell]$. After scanning each w_i , every entry in w has an associated vertex label, which we hereafter just refer to as vertices.
- **Step three:** In this step we construct a circular list of four oriented edges for each vertex. Let w'_i denote the string w_i , but with each entry $w_i[k]$ replaced with its associated vertex. Assume that curves i and j intersect at vertex v_1 as is shown below

$$\begin{aligned} w'_i: & \cdots v_0 v_1 v_2 \cdots, \\ w'_j: & \cdots v_3 v_1 v_4 \cdots. \end{aligned}$$

By computing the parity of the number of intersections between i and j along w_i , we can determine whether curve i is moving into the interior of curve j at v_1 , or whether it is moving into the exterior. Since each w_k is being traversed in a clockwise direction starting at the unbounded face, the interior of w_k is always to its right. Thus, if the parity is odd, then we are moving into the interior, and if the parity is even, then we are moving into the exterior. In the former case the circular order of oriented edges incident to v_1 must be

$$\{(v_1, v_0)_i, (v_1, v_4)_j, (v_1, v_2)_i, (v_1, v_3)_j\},$$

and in the latter case, the circular order must be

$$\{(v_1, v_0)_i, (v_1, v_3)_j, (v_1, v_2)_i, (v_1, v_4)_j\}.$$

The notation $(v_1, v_0)_i$ indicates the oriented edge from v_1 to v_0 along curve i . Using the same method for each vertex, we get a rotation system that uniquely identifies a 2-cell embedding of the Venn diagram, up to isomorphism.

The rotation system only depends on which string of w , x , y , or z is chosen as the first string of Grünbaum encoding. Since there exist permutation mappings to deduce all other strings from any of w , x , y , or z , all rotation systems that arise from the three steps are equivalent up to isomorphism. Therefore, the Grünbaum encoding uniquely identifies the Venn diagram. \square

To end this section we note that the Grünbaum encoding can be used to determine whether a Venn diagram is polar symmetric or whether it has antipodal symmetry (in both instances, given that the north and south poles have been fixed). The diagram is polar symmetric if there are integers k and k' such that $w_i = z_{i+k}$ and $x_i = y_{i+k'}$, where index computations are taken mod n . Similarly, the diagram has antipodal symmetry if there are integers k and k' such that $w_i = x_{i+k}$ and $y_i = z_{i+k'}$, where index computations are taken mod n .

3.2. The matrix representation

Every simple monotone n -Venn diagram is exposed and thus every 1-region is adjacent to the empty region. So the empty region surrounds a “ring” of $\binom{n}{1}$ 1-regions. An intersection point is said to be part of ring i if of the four incident regions, two are in ring i and the other two are in rings $i - 1$ and $i + 1$. Since each region is started by one intersection point and ended by another one, there are $\binom{n}{1}$ intersection points in the first ring. Similarly, every 2-region is adjacent to at least one 1-region. So, there are $\binom{n}{2}$ 2-regions that form the second ring surrounded by the first ring and which contains $\binom{n}{2}$ intersection points. In general, in a simple monotone n -Venn diagram, there are $n - 1$ rings of regions, where all regions in a ring are in the interior of the same number of curves and every region in ring i , $1 \leq i \leq n - 1$, is adjacent to at least one region in ring $i - 1$ and also to at least one region in ring $i + 1$. The number of intersection points in the i th ring is the same as the number of regions in the i th ring, which is $\binom{n}{i}$. The rings have different colours in Fig. 1(c).




Fig. 5. Matrix representation of one slice of the symmetric 7-Venn diagram of Fig. 1(c), the diagram known as “Adelaide.”

Thus a simple monotone n -Venn diagram can be represented by an $(n-1) \times (2^n - 2)$ binary matrix with exactly one 1 in each column, where each 1 in the matrix represents an intersection point in the Venn diagram. Row i of the matrix corresponds to ring i of the Venn diagram, which will thus contain $\binom{n}{i}$ 1's. Furthermore, because there are no 2-faces, we may assume that there are no two adjacent 1s in any row. Fig. 5 shows the matrix representation of $1/7$ -th of the symmetric 7-Venn diagram of Fig. 1(c).

Because of the property of symmetry, a symmetric n -Venn diagram can be partitioned to n identical circular sectors, where each sector is specified by two rays from the point of symmetry offset by a central angle of $2\pi/n$ radians. Therefore, having an n th of the matrix, which we call the *slice matrix*, is enough to represent the diagram. We simply copy the slice matrix n times to get the entire matrix representation.

If matrix M is a representation of a Venn diagram, then any matrix obtained by a circular shift of M by some number of columns is also a representation of the same Venn diagram. Therefore, we can always shift M such that the first entry of the first column is a 1. Such a slice matrix with a 1 bit at the upper left entry is called a *standard matrix*.

Note that there are $\binom{m-k+1}{k}$ binary sequences of length m that have k 1s, with the restriction that no 1s are adjacent. Thus we can generate all possible standard matrices by generating the n distinct combinations, whose sizes, in the case of $n=7$, correspond to the following product of binomial coefficients:

$$\binom{18}{0} \cdot \binom{17}{3} \cdot \binom{14}{5} \cdot \binom{9}{5} \cdot \binom{4}{3} \cdot \binom{1}{1} = 686\,125\,440. \quad (1)$$

This is a large but manageable number of possibilities.

Given a standard matrix representation M , we need to check that there are no two regions with the same rank to check if it represents a valid Venn diagram. To compute the rank of a region, we need to know the label of the curves that contain the region. For this purpose it is useful to consider another matrix, which we call the P -matrix. It consists of n rows and $2^n - 2$ columns. The initial column is the identity, and each successive column is obtained from its predecessor by swapping the two values in a column if there is a 1 in the corresponding column of matrix M^n (M copied n times), that is, if $M_{i,j}^n = 1$ then $P_{i,j+1} = P_{i+1,j}$ and $P_{i+1,j+1} = P_{i,j}$. For example the first 19 columns of the P -matrix for the matrix M of Fig. 5 is given below.

0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
1	0	2	2	2	4	4	4	4	4	3	3	3	3	3	3	3	3
2	2	0	3	3	4	2	2	5	5	3	4	4	4	0	0	0	0
3	3	3	0	4	3	3	5	2	3	5	5	5	0	4	5	5	5
4	4	4	4	0	0	5	3	3	2	2	2	0	5	4	2	2	2
5	5	5	5	5	5	0	0	0	0	0	0	2	2	2	4	6	6
6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	4	4

Suppose that the vector $[c_0, c_1, \dots, c_{n-1}]$ is a column of the P -matrix, where c_0 is the label of the outermost (top) curve and c_{n-1} is the label of the innermost (bottom) curve. Then the region at ring i , $1 \leq i \leq n-1$, is in the interior of curves c_0, \dots, c_{i-1} and the rank of the region is $\sum_{k=0}^{i-1} 2^{c_k}$. If the next column of P is obtained by exchanging c_i and c_{i+1} , then we have created exactly one new region, a region with rank $\sum_{k=0}^i 2^{c_k} - 2^{c_{i-1}}$. Proceeding from left to right through P , we can then compute the rank of each region. Matrix P (and hence matrix M) represents a valid simple monotone Venn diagram





Fig. 6. An i -face with p lower vertices.



$$\begin{aligned}
 & ((3, 3, 3, 3, 2, 1), \\
 & (1, 1, 2, 1, 1, 2, 1, 1, 2, 1, 1, 2), \\
 & (1, 0, 1, 1, 1, 0, 1, 1, 0, 1, 1, 0, 1, 1, 0, 1, 1), \\
 & (1, 1, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0)) \\
 & \quad (b)
 \end{aligned}$$

Fig. 7. (a) Cylindrical representation of 6-Venn diagram, (b) composition representation of (a).

if we get exactly $2^n - 2$ regions with distinct ranks and the identity permutation $C = [0, 1, \dots, n - 1]$ after applying the final exchange. Initially, the first column is the identity also, and the regions corresponding to that first column have ranks $1, 3, 7, \dots, 2^{n-1} - 1$.

3.3. Composition representation

In this subsection we introduce a representation of Venn diagrams that is based on integer compositions. In this representation, we use a sequence of non-negative integers to show the size of faces in each ring and also to specify the position of intersection points of the next ring relative to the position of intersection points in the current ring. A k -composition of n is a sequence of non-negative integers (a_1, a_2, \dots, a_k) such that $n = \sum a_i$. Let $\mathcal{C}(n, k)$ denote the set of all k -compositions of n .

In a simple monotone n -Venn diagram there are $\binom{n}{i+1}$ intersection points at ring $i + 1$ that are distributed among $\binom{n}{i}$ intersection points at ring i . Starting from an intersection point at ring i , we can specify the number of intersection points at ring $i + 1$ between each two consecutive intersection points of ring i using a composition of $\binom{n}{i+1}$ into $\binom{n}{i}$ parts. Consider a generic face in ring i , like that shown in Fig. 6. It is delimited by some two intersection points that are adjacent to regions on the same ring. If there are p intersections on the lower part of the face (such intersections are adjacent to k -faces with $k \geq i$), then p will be part of the composition for ring i .

A simple monotone n -Venn diagram V can be represented by a sequence $C_V = (c_1, c_2, \dots, c_{n-2})$ where

$$c_i \in \mathcal{C}\left(\binom{n}{i+1}, \binom{n}{i}\right).$$

The composition c_i is determined by following ring i in a circular fashion, recording for each i -face the number of intersection points that only lead to k faces, where $k > i$. Thus the underlying compositions are circular. To be able to recover the diagram from the compositions, we need to specify where each composition starts. Given $c_i = p + q + r + \dots$, the starting face for c_{i+1} is one that lies below the face corresponding to p , call it F . Suppose that F is joined to the rest of the i th ring by vertices u and v on its left and right, respectively. If $p > 1$, then the face is the one that lies between the first two of the lower p intersection points (see Fig. 6). If $p = 1$, then the face is the one that lies between the lower vertex and v . If $p = 0$, then the face is the one that lies below the edge from u to v . With these conventions, we say that C_V is a *composition representation* of V . A simple monotone Venn diagram does not necessarily have a unique composition representation because starting from a different intersection point on the first ring we may get a different composition representation of the Venn diagram. Fig. 7 shows the composition representation of the 6-Venn diagram of Fig. 2. The black dots indicate where the various faces of the compositions start.

We now list several observations that will help us cut down the size of the search space of the generating algorithm.




Fig. 8. Non-adjacent 1's in the first ring composition.

Remark 1. For any simple monotone n -Venn diagram V , the largest part of c_i in the composition representation is at most $n - i - 1$.

Proof. A region at ring i is in the interior of i curves. Since the size of a region is at most n and no two edges belong to the same curve by Lemma 2, at most $n - i$ remaining curves can be used to shape the region. As shown in Fig. 6, to put p intersection points between the two end points of the region on the next ring, we need $p + 1$ curves, $p - 1$ curves for the bottom side and two curves for the left and right sides. So, $p \leq n - i - 1$. \square

Remark 2. In the composition representation of any simple monotone Venn diagram with more than 3 curves, there are no two non-adjacent 1's in c_1 .

Proof. Suppose, there is such a Venn diagram V , then the first ring of the Venn diagram will be like Fig. 8, where regions A and D correspond to non-adjacent 1's in the composition and $A \neq D$. Then $A \cap D = \emptyset$ which contradicts the assumption that V is a Venn diagram. So in the first ring composition there are at most two 1's which must be adjacent. \square

Remark 3. There are no two consecutive 0's in c_2 for the composition representation of any simple monotone n -Venn diagram.

Proof. By Lemma 4. \square

Definition 1. Let $r_1, r_2 \in \mathcal{C}(n, k)$ be two compositions of n into k parts. r_1 and r_2 are *rotationally distinct* if it is not possible to get r_2 from any rotation of r_1 or its reversal.

Let \mathcal{F}_n denote the set of all rotationally distinct compositions of $\binom{n}{2}$ into n parts such that for any $r \in \mathcal{F}_n$ there are no two non-adjacent parts of 1 and all parts are less than or equal to $n - 2$.

Lemma 8. If c_1 is the composition corresponding to the first ring of a simple monotone n -Venn diagram, then $c_1 \in \mathcal{F}_n$.

Proof. Given a simple monotone n -Venn diagram, suppose we get the composition representation P of V by picking a particular intersection point x in the first ring as the reference point. Now let P' be another representation of V using any other intersection point different than x as the reference point. It is clear that c'_1 in P' is a rotation of c_1 in P . Also for any composition representation P'' of the mirror of V the first composition c''_1 in P'' is a rotation of the reversal of c_1 . By the Remarks 1 and 2 the largest part of c_1 is $n - 2$ and there are no two non-adjacent 1's in c_1 . Therefore, there is a composition $c \in \mathcal{F}_n$ which is rotationally identical to c_1 . \square

3.4. Permutation sequence

Recall that when testing whether a standard matrix represented a Venn diagram we used the P-matrix. But instead of storing the P-matrix as a sequence of permutation, we could simply record the row in which the intersection occurs, as is shown in Fig. 9. If π and π' are two successive permutations then the corresponding entry in the permutation sequence is i if π' is obtained from π by exchanging $\pi(i)$ and $\pi(i - 1)$. We call the resulting sequence of length $2^n - 2$ the *permutation sequence*. It is clear that there are $\binom{n}{i}$ elements of value i in the sequence, where $i \in \{1, 2, \dots, n - 1\}$. Furthermore, since we assume that the first permutation of the P-matrix is the identity, the first entry of the permutation sequence is a 1.

Using the matrix representation or the composition representation we generate Venn diagrams vertically (top-to-bottom). However, permutation sequences generate Venn diagrams horizontally from left-to-right. This allows us to immediately compute the rank of regions as we move from one permutation to the next, and thus can potentially help in reducing the size of the backtracking tree, since non-Venn diagrams are recognized early.




Fig. 9. The P -matrix of one slice of Grünbaum's 5-ellipses Venn diagram. Shown below it are the first six values of the permutation sequence.

4. Generating algorithms

4.1. Generating simple symmetric convex 7-Venn diagrams

We used the matrix representation to generate all simple symmetric monotone 7-Venn diagrams. The generating algorithm is shown in Algorithm 1. A slice matrix of a simple symmetric monotone 7-Venn diagram has 6 rows and 18 columns where the number of 1's in rows 1, 2, ..., 6 are 1, 3, 5, 5, 3, 1 respectively. To generate each row we are generating restricted combinations; e.g., all bitstrings of length 18 with k 1's, no two of which are adjacent. The total number of combinations generated is the number given in (1).

For each generated slice matrix, we copy it n times to get a complete matrix M , which we then convert into the corresponding P -matrix. Then we check if P represents a valid Venn diagram by counting the number of distinct regions. This process was described earlier in the paper.

To eliminate isomorphic Venn diagrams, we use the Grünbaum encoding, which is easily created from the P -matrix in $O(2^n)$ time. Since the diagram is symmetric, we only need w_i for $i = 0$ and similarly, we only need x_0 , y_0 , and z_0 (which are easily obtained from w_0 using the permutation mappings explained earlier). We first initialize a vector ρ with the permutation of $\{0, 1, \dots, n-1\}$ such that the curves appear on the outer boundary of the outermost ring in the order $\rho(1), \rho(2), \dots, \rho(n-1)$. To get w_0 , we scan P left-to-right and follow curve 0's intersections with other curves, translating them by ρ^{-1} . The lexicographically smallest string of $\{w_0, x_0, y_0, z_0\}$ is chosen as the representative Grünbaum encoding of the Venn diagram. Comparing the Grünbaum encodings of the previously generated Venn diagrams with the Grünbaum encoding of the current Venn diagram, we can eliminate isomorphic Venn diagrams.

Algorithm 1 Generating all simple symmetric convex 7-Venn diagrams

```

1: procedure GENSYMSEVEN( $i$ )
2:    $G \leftarrow \{\}$ 
3:    $V \leftarrow \{\}$ 
4:   for each standard slice matrix  $M$  do
5:     for  $i \leftarrow 1, \dots, 7$  do
6:        $X \leftarrow X \cdot M$ 
7:     end for
8:     if isVenn( $X$ ) then
9:        $g \leftarrow$  Grünbaum encoding of  $X$ 
10:      if  $g \notin G$  then
11:         $G \leftarrow G \cup g$ 
12:         $V \leftarrow V \cup X$ 
13:      end if
14:    end if
15:  end for
16: end procedure

```

Using this algorithm we found exactly 23 simple symmetric monotone 7-Venn diagrams of which 6 are polar symmetric. See Figs. 10 and 11 for attractive renderings of these diagrams.

These computations were checked by using an algorithm based on the composition representation, and using a depth-first-search labeling algorithm for the isomorphism check.

4.2. Generating simple polar symmetric 6-Venn diagrams

Given the upper/lower half of the cylindrical representation of a polar symmetric Venn diagram, one can generate the whole diagram by creating a copy of the given half, turning it upside down and rotating it until the two parts match together. So, to generate a polar symmetric monotone Venn diagram, we need only to generate the first $\lceil \frac{n-2}{2} \rceil$ compositions.

Two halves of the diagram can match only if gluing them using the intersection points doesn't create any faces of size 2. Given the last composition of the upper half, for each positive part a_j there are $a_j - 1$ edges that bound the corresponding




Fig. 10. All simple monotone polar symmetric 7-Venn diagrams, using the names given to them by Edwards [7]. Around each diagram is the lexicographically smallest Grünbaum encoding.

face from the bottom and there is a gap between two faces corresponding to two consecutive parts of the composition. So, we can map the composition to a bit-string where each 1 represents a bounding edge of a face and each 0 represents the gap between two faces. The length of bit-string is the same as the sum of all parts of the composition. In other words the composition (a_1, a_2, \dots, a_k) is mapped to the following bit-string.

$$\underbrace{11 \dots 1}_{a_1-1 \text{ bits}} \underbrace{011 \dots 1}_{a_2-1 \text{ bits}} \dots \underbrace{011 \dots 1}_{a_k-1 \text{ bits}} 0$$

We can find all matchings of the two halves by computing the bitwise “and” of the bit-string and its reverse for all left rotations of the reverse bit-string. Any result other than 0 means that there is at least one face of size 2 in the middle. Then for each matching we compute the matrix representation of the resulting diagram. The matrix can be obtained by sweeping the compositions from left to right and computing the position of each intersection point. Checking each resulting matrix for all compositions gives us all possible polar symmetric 6-Venn diagrams.

Algorithm 2 Generating all simple convex polar-symmetric 6-Venn diagrams

```

1: procedure GENPOLARSIX( $i$ )
2:   for each composition  $(a_1, a_2, \dots, a_6) \in \mathcal{F}_6$  do
3:     for each composition  $(b_1, b_2, \dots, b_{15}) \in \mathcal{C}(20, 15)$  do
4:       create the corresponding upper and lower halves
5:       for  $i \leftarrow 1, \dots, 20$  do
6:         glue the upper and lower halves
7:         if there are no parallel edges in the diagram then
8:           compute matrix  $X$  representing the diagram
9:           if  $\text{isVenn}(X)$  then
10:            add  $X$  to the list of Venn diagrams
11:           end if
12:         end if
13:         rotate lower half one point to the left
14:       end for
15:     end for
16:   end for
17: end procedure

```

Using the exhaustive search based on this algorithm we found 375 simple monotone polar symmetric 6-Venn diagrams. This result was independently checked by using a separate program that is based on a different search method and which is described in the next subsection. Table 2 shows the number of Venn diagrams for each particular composition of the first level.




Fig. 11. All 17 simple monotone symmetric 7-Venn diagrams that do not have polar symmetry. Around each diagram is the lexicographically smallest Grünbaum encoding.

Table 2

Number of polar symmetric 6-Venn diagrams for \mathcal{F}_6 .

Composition	Venn diagrams	Composition	Venn diagrams
4 4 3 2 1 1	25	2 3 4 3 2 1	5
4 3 4 2 1 1	0	4 2 2 4 2 1	0
3 4 4 2 1 1	38	3 3 2 4 2 1	9
4 4 2 3 1 1	0	2 4 2 4 2 1	0
4 3 3 3 1 1	12	3 2 3 4 2 1	3
3 4 3 3 1 1	9	4 3 2 2 3 1	9
4 2 4 3 1 1	0	3 4 2 2 3 1	15
4 3 2 4 1 1	0	4 2 3 2 3 1	0
4 4 2 2 2 1	15	3 3 3 2 3 1	9
4 3 3 2 2 1	2	3 2 4 2 3 1	0
3 4 3 2 2 1	30	4 2 2 3 3 1	12
4 2 4 2 2 1	0	3 3 2 3 3 1	4
3 3 4 2 2 1	6	4 2 2 2 4 1	0
2 4 4 2 2 1	13	4 3 2 2 2 2	15
4 3 2 3 2 1	7	4 2 3 2 2 2	22
3 4 2 3 2 1	8	3 3 3 2 2 2	6
4 2 3 3 2 1	6	4 2 2 3 2 2	1
3 3 3 3 2 1	41	3 3 2 3 2 2	21
2 4 3 3 2 1	22	3 2 3 2 3 2	3
3 2 4 3 2 1	7		

4.3. Generating simple convex 6-Venn diagrams

In this subsection we explain the algorithm for generating all simple convex 6-Venn diagrams using what we call the *permutation representation*. The algorithm is based on a simple idea: Starting from the identity permutation, we generate via a backtracking program, a sequence of permutations. For each permutation we try all possible exchanges of adjacent curves to get the next permutation until we get a sequence of length 62 that represents a Venn diagram. The height of the recursion tree of the backtracking algorithm is thus 62. There are five possible choices for the first exchange and four choices for each of the other exchanges. So, there are $5 \cdot 4^{61} \approx 2.66 \cdot 10^{37}$ possible sequences in total. Therefore, we need some good rules to prune the search tree as much as possible.

As with matrix representation, if a sequence of exchanges X represents a valid Venn diagrams then any rotation of X represents the same Venn diagram. A sequence of exchanges X is *canonic* if among all rotation of X , it has the largest corresponding sequence of permutations. We use the canonical form to eliminate all sequences which are identical to a canonic sequence up to rotations. Given a prefix of length k of an exchange sequence $S = (s_1 s_2 \dots s_{62})$, if there is some i , with $1 < i \leq k$, such that starting at position i in S with identity permutation we get a larger permutation sequence then S is not canonic. So we can check the canonicity for each generated prefix of exchanges and reject the non-canonic exchange sequences as soon as possible.

As another rule, if there are two exchanges in two adjacent permutations such that the positions of the exchanges do not overlap, then the exchange of the lower position comes first. Because the resulting diagram in both cases is the same and we don't need to generate both of them.

As was mentioned before, we can compute and check the rank of regions as we move from one permutation to the next. So we can cut a non-Venn diagram at the earliest stage of recursion. Also, not explicit in the pseudocode on the next page, in our actual code we also exploit some other simple properties of convex Venn diagrams, such as [Lemma 2](#), to speed up the program.

The pseudocode for generating all simple convex 6-Venn diagram is shown in [Algorithm 3](#). Input i is the next exchange. S is the exchange sequence and V is the list of Venn diagrams that have been found so far. The current permutation of curve labels is stored in vector C and vector $rank$ is used to store the rank of current region of each ring. The number of regions that have been visited so far is stored in rno and vector $visited$ is used to keep track of the visited regions. We start with the identity permutation as the curve labels. So, vector $rank$ must be initialized to $[1, 3, 7, 15, 31]$, because as was mentioned before, the rank of the region at ring i is $\sum_{k=0}^{i-1} 2^{C_k}$. Two consecutive regions at ring i only differ in $C[i]$ and $C[i+1]$. To update the rank vector, after swapping $C[i]$ and $C[i+1]$, we need to add curve $C[i]$ to the new region and exclude $C[i+1]$ from it. The rank of regions of the other rings remain unchanged. Lines 24 to 27 of the algorithm restore the variables to their state before the recursive call.

Algorithm 3 Generating all simple convex 6-Venn diagrams

```

1: procedure GENSIX( $i, rno$ )
2:   if  $S$  is not canonic then
3:     return
4:   end if
5:   if  $rno = 62$  then
6:     if  $C = id$  then
7:        $V \leftarrow V \cup S$ 
8:     end if
9:     return
10:  end if
11:   $r \leftarrow rank[i]$ 
12:  if  $visited[r] = 1$  then
13:    return
14:  end if
15:   $visited[r] \leftarrow 1$ 
16:   $S \leftarrow S \cup \{i\}$ 
17:   $C[i] := C[i+1]$ 
18:   $rank[i] \leftarrow rank[i] + 2^{C[i]} - 2^{C[i+1]}$ 
19:  for  $j \leftarrow i-1, \dots, n-1$  do
20:    if  $j \neq i$  then
21:      GENSIX( $j, rno+1$ )
22:    end if
23:  end for
24:   $rank[i] \leftarrow rank[i] - 2^{C[i]} + 2^{C[i+1]}$ 
25:   $C[i] := C[i+1]$ 
26:   $S \leftarrow S \setminus \{i\}$ 
27:   $visited[r] \leftarrow 0$ 
28: end procedure

```

A distributed version of the algorithm takes only a few hours on a machine with 64 processors to generate all simple monotone 6-Venn diagrams. There are 39 020 such Venn diagrams in total.




Fig. 12. All simple convex spherical 6-Venn diagrams whose isometry groups have order 8. The diagrams on each row can be transformed, one into the other, by flipping an opposing pair of quadrants.

4.4. Testing for symmetry


In this subsection we describe how we tested each of the 39 020 planar monotone 6-Venn diagrams to determine whether they had any non-trivial automorphism when embedded on the sphere.

It is well-known that a 3-connected planar graph has a unique embedding (under the assumption that reversing the sense of clockwise for an embedding gives an equivalent embedding). Because of this property, 3-connected planar graphs can be put into a canonical form, and it is possible to compute the automorphism group using a very simple algorithm based on a special type of breadth-first search (called a *clockwise BFS*) that runs in $O(n^2)$ time in the worst case. It is not clear who originally came up with this elegant algorithm. One place it has been explained and used is [8].

Given a rotation system for a graph, a *clockwise breadth-first search* (BFS) starts at a specified root vertex r , and has a specified neighbor f of r designated to be the first child of r . A BFS is performed with the restrictions that the neighbors of r are traversed in clockwise order starting with the first child f . When the neighbors of a non-root vertex u are visited, they are traversed in clockwise order starting with the BFS parent of u . A clockwise BFS labels each vertex with its breadth-first index.

To get the canonical form for an embedding, consider all possible selections of a root vertex, a first child, and the direction representing clockwise and choose one giving a lexicographically minimized rotation system. The selections giving an identically labeled rotation system give the automorphisms. If there are no automorphisms which have different choices for the clockwise direction, then the embedding is said to be *chiral*, and otherwise it is *achiral*.

The algorithm above gives the automorphisms that can be realized when embedding the graph on the sphere (these will be called the *spherical automorphisms*). To consider only those that map the innermost face (i.e., the face that is interior to all the curves) to itself, one trick that can be used is to embed a new vertex w inside the innermost face and connect it to



$x_0 = y_0$	1 2 4 3 4 2 1 4 1 5 3 1 2 1 4 2 3 2 4 1 2 4 3 6 1 2 4 3 4 2 1 4 1 5 3 1 2 1 4 2 3 2 4 1 2 4 3 6
$x_1 = y_1$	2 4 3 4 2 0 4 0 3 5 0 2 0 4 2 3 2 4 0 6 3 2 4 0 2 4 3 4 2 0 4 0 3 5 0 2 0 4 2 3 2 4 0 6 3 2 4 0
$x_2 = y_2$	3 1 0 4 5 3 1 0 4 1 0 3 0 1 4 0 4 3 6 4 1 4 0 1 3 1 0 4 5 3 1 0 4 1 0 3 0 1 4 0 4 3 6 4 1 4 0 1
$x_3 = y_3$	4 1 5 0 2 4 2 0 1 0 2 6 4 1 0 1 4 2 4 1 5 0 2 4 2 0 1 0 2 6 4 1 0 1 4 2
$x_4 = y_4$	5 2 1 0 3 0 1 2 0 2 6 3 2 1 2 0 1 3 1 0 2 1 0 3 5 2 1 0 3 0 1 2 0 2 6 3 2 1 2 0 1 3 1 0 2 1 0 3
$x_5 = y_5$	6 4 1 3 0 2 6 2 0 3 1 4
$x_6 = y_6$	0 2 3 4 1 5 1 4 3 2 0 5

Fig. 13. A spherical 7-Venn diagram with a 4-fold rotational symmetry and the polar symmetry. Shown below the diagram is its lexicographically smallest Grünbaum encoding.

all the vertices on the innermost face. Then apply the clockwise BFS but only with the selections that have w as the root vertex. These are the *planar automorphisms* (they are the automorphisms realizable for a picture of an embedding drawn on the plane).

The *polar automorphisms* are those which map the innermost face either to itself or to the external face. The trick used to compute these is to add two new vertices; w_1 which is connected to the vertices on the innermost face, and w_2 which is connected to the vertices on the external face. Then apply the clockwise BFS selections that have either w_1 or w_2 as the root vertex.

Using the above described ideas we found that 375 of them are polar symmetric, which confirms our results of generating polar symmetric 6-Venn diagrams using compositions. And 270 of the generated Venn diagrams have the (antipodal) rotary reflection symmetry. Amongst all simple convex 6-Venn diagrams, we found 27 Venn diagrams with automorphism group order of 4. There are only 6 Venn diagrams that have an automorphism group order of 8; these are shown in Fig. 12. In all six cases, there are two curves that intersect in only two points. Imagine each Venn diagram as being projected on a sphere such that these two curves map to two great circles that meet perpendicularly at the north and south poles. The three cases on the left side of Fig. 12 have the 4-fold rotational symmetry about an axis through the poles. The three other cases have the reflection symmetry across the two planes which pass through the two great circles. The other factor of 2 comes from the fact that all of these diagrams are polar-symmetric.

5. Concluding remarks

In this paper we investigated general techniques for generating simple monotone Venn diagrams. We proved that every Grünbaum encoding uniquely identifies a simple monotone Venn diagram up to isomorphism and we extended the proof to simple exposed Venn diagrams where every curve touches the empty region. We introduced different representation of simple monotone Venn diagrams and we described some generating algorithm based on these representations to exhaustively list Venn diagrams. Using these algorithm we showed that there are 23 simple monotone symmetric 7-Venn diagrams

where amongst of them six cases are polar-symmetric. For simple monotone Venn diagrams on six curves, we showed that there are 39020 diagrams in total, 375 of which are polar-symmetric.

It remains to do an exhaustive enumeration of all monotone simple 7-Venn diagrams. Prior to this paper, the only spherical symmetries known were 7-fold rotations and polar flips. However, by hand we have created a spherical 7-Venn diagram with a 4-fold rotational symmetry and a polar symmetry; see Fig. 13. The lexicographically smallest Grünbaum encoding of the diagram is shown below it. Note that because of the polar symmetry x and y (also w and z) strings of the Grünbaum encoding are identical. Undoubtedly there are other such examples of 7-Venn diagrams with isometry groups of order 4 and 8.

Acknowledgements

We thank the referees for their insightful comments which led to improvements in the paper. At one of the referees suggestions we have made some of our software available. See the page <http://webhome.cs.uvic.ca/~ruskey/Publications/SixVenn/SixSevenVenn.html>.

References

- [1] G. Brinkmann, B.D. McKay, Fast generation of planar graphs, *MATCH Commun. Math. Comput. Chem.* 58 (2007) 323–357.
- [2] B. Bultena, B. Grünbaum, F. Ruskey, Convex drawings of intersecting families of simple closed curves, in: 11th Canadian Conference on Computational Geometry, pp. 18–21.
- [3] T. Cao, Computing all the simple symmetric monotone Venn diagrams on seven curves, Master's thesis, University of Victoria, Canada, 2001.
- [4] T. Cao, K. Mamakani, F. Ruskey, Symmetric monotone Venn diagrams with seven curves, in: P. Boldi, L. Gargano (Eds.), *Fun with Algorithms*, in: *Lecture Notes in Computer Science*, vol. 6099, Springer, Berlin/Heidelberg, 2010, pp. 331–342.
- [5] K.B. Chilakamarri, P. Hamburger, R.E. Pippert, Venn diagrams and planar graphs, *Geom. Dedicata* 62 (1996) 73–91.
- [6] J.R. Edmonds, A combinatorial representation of polyhedral surfaces, *Notices Amer. Math. Soc.* 7 (1960).
- [7] A.W.F. Edwards, Seven-set Venn diagrams with rotational and polar symmetry, *Combin. Probab. Comput.* 7 (1998) 149–152.
- [8] P. Fowler, D. Horspool, W. Myrvold, Vertex spirals in fullerenes and their implications for nomenclature of fullerene derivatives, *Chem. Eur. J.* 13 (2007) 2208–2217.
- [9] J. Griggs, C.E. Killian, C.D. Savage, Venn diagrams and symmetric chain decompositions in the Boolean lattice, *Electron. J. Combin.* 11 (2004), Research Paper 2, 30 pp. (electronic).
- [10] B. Grünbaum, Venn diagrams and independent families of sets, *Math. Mag.* 48 (1975) 12–23.
- [11] B. Grünbaum, Venn diagrams. II, *Geombinatorics* 2 (1992) 25–32.
- [12] D.W. Henderson, Venn diagrams for more than four classes, *Amer. Math. Monthly* 70 (1963) 424–426.
- [13] K. Mamakani, W. Myrvold, F. Ruskey, Generating all simple convexly-drawable polar symmetric 6-Venn diagrams, in: C.S. Iliopoulos, W.F. Smyth (Eds.), *Combinatorial Algorithms*, in: *Lecture Notes in Computer Science*, vol. 7056, Springer, Berlin/Heidelberg, 2011, pp. 275–286.
- [14] F. Ruskey, C.D. Savage, S. Wagon, The search for simple symmetric Venn diagrams, *Notices Amer. Math. Soc.* 53 (2006) 1304–1312.
- [15] F. Ruskey, M. Weston, A survey of Venn diagrams, *Electron. J. Combin.* 11 (2004) 22.
- [16] F. Ruskey, M. Weston, Spherical venn diagrams with involutory isometries, *Electron. J. Combin.* 18 (2011), paper #P191, 14 pp.
- [17] S. Wagon, P. Webb, Venn symmetry and prime numbers: A seductive proof revisited, *Amer. Math. Monthly* 115 (2008) 645–648.
- [18] M. Weston, Symmetries of Venn diagrams on the sphere, Ph.D. thesis, University of Victoria, Canada, 2009.
- [19] H. Whitney, A theorem on graphs, *Ann. of Math.* (2) 32 (1931) 378–390.