Face-balanced, Venn and polyVenn diagrams

by

Bette Bultena
B.Sc., University of Victoria, 1995
M.Sc., University of Victoria, 1998

A Dissertation Submitted in Partial Fulfillment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Computer Science

Supervisory Committee

Dr. Frank Ruskey, Supervisor
(Department Computer Science)

Dr. Wendy Myrvold, Departmental Member
(Department of Computer Science)

Dr. Dale Olesky, Departmental Member
(Department of Computer Science)

Dr. Anthony Quas, Outside Member
(Department of Mathematics)

Supervisory Committee

Dr. Frank Ruskey, Supervisor
(Department Computer Science)

Dr. Wendy Myrvold, Departmental Member
(Department of Computer Science)

Dr. Dale Olesky, Departmental Member
(Department of Computer Science)

Dr. Anthony Quas, Outside Member
(Department of Mathematics)

ABSTRACT

A *simple n-Venn diagram* is a collection of n simple intersecting closed curves in the plane where exactly two curves meet at any intersection point; the curves divide the plane into 2^n distinct open regions, each defined by its intersection of the interior or exterior of each of the curves. A Venn diagram is *reducible* if there is a curve that, when removed, leaves a Venn diagram with one less curve and *irreducible* if no such curve exists. A Venn diagram is *extendible* if another curve can be added, producing a Venn diagram with one more curve. Currently it is not known whether every simple Venn diagram is extendible by the addition of another curve. We show that all simple Venn diagrams with 5 curves or less are extendible to another simple Venn diagram. We also show that for certain Venn diagrams, a new extending curve is relatively easy to produce.

We define a new type of diagram of simple closed curves where each curve divides the plane into an equal number of regions; we call such a diagram a *face-balanced*

diagram. We generate and exhibit all face-balanced diagrams up to and including those with 32 regions; these include all the Venn diagrams.

Venn diagrams exist where the curves are the perimeters of polyominoes drawn on the integer lattice. When each of the 2^n intersection regions is a single unit square, we call these *minimum area polyomino Venn diagrams*, or *polyVenns*. We show that polyVenns can be constructed and confined in bounding rectangles of size $2^r \times 2^c$ whenever $r, c \geq 2$ and $n = r + c$. We show this using two constructive proofs that extend existing diagrams. Finally, for even n, we construct polyVenns with n polyominoes in $(2^{n/2} - 1) \times (2^{n/2} + 1)$ bounding rectangles in which the empty set is not represented as a unit square.

Contents

List of Tables

List of Figures

ACKNOWLEDGEMENTS

Once again, I am very grateful to have the privilege and pleasure to research such beautiful and fascinating mathematical diagrams. It would not have been possible without the generosity and brilliance of people who spend their lives in discovery and then gladly share their successes and failures with the rest of us. For me, the best of these is my supervisor, Frank Ruskey, who always inspires me with his humble brilliance. Thank-you, Frank! I would also like to express my gratitude to the amazing Branko Grünbaum, for his very kind encouragement and guidance. I am very fortunate to have such mentors.

Many thanks also to the National Science and Engineering Research Council of Canada for investing in my education and making it possible for me to take a leave from work to concentrate on my studies.

I thank my husband, Richard Baldwin, who helped me up, dusted me off, and showed me the path whenever I became discouraged. During times of inspiration, he stayed up with me to share both good ideas and those that needed some correction. Lastly, but by no means least, I thank my sons, Adrian and Gabe Letourneau, for being so supportive of me as we all struggled with our schoolwork together.

Fifty percent of success is in just showing up.
Woody Allen

DEDICATION

To my parents, now deceased, who were always proud of me. Dear Dad, as promised, your names are in my thesis.

To Pieta and Siewert Bultena, who continue to live inside me and sometimes borrow my eyes so they can see the world they both loved so much.

Chapter 1

Introduction

In spite of the continued interest in Venn diagrams, an almost 30 year old conjecture remains an open problem [25]. It is a surprise that it has not been proven; it seems such a simple thing: Is it possible to extend a Venn diagram of n curves to one of $n + 1$ curves? John Venn [23] showed that it is true if the Venn diagram is itself *reducible*, meaning that the removal of a particular curve results in a Venn diagram. Kiran B. Chilakamarri et al. [8] proved that it is true if more than two curves can cross at a point, but the problem remains open for simple Venn diagrams, where only two curves cross at a point.

In this thesis, the first problem we tackle is what makes a Venn diagram extendible. The answer to solving this problem may lie in how we define Venn diagrams. Venn diagrams are generally thought of as a set of Jordan curves in the plane, or as planar graphs. We show that Venn diagrams are a subset of a new and interesting set of curves. We show that they can also be rendered as a collection of polyominoes on a grid.

Whether we look at Venn diagrams as graphs, as collections of curves, polyominoes or as a property that acts on sets of objects, we ask the same questions:

- What do they look like and can we draw them?
- How many are there for n curves or polyominoes?
- Are they reducible?
- Are they extendible?

We give some relevant graph theory definitions, particularly relating to planar graphs, in the remainder of this chapter. In Chapter 2, we define sets of diagrams

consisting of connected simple Jordan curves on the plane. We also generate and count all such diagrams that can be embedded on the sphere. In Chapter 3, we introduce a newly discovered set of diagrams where every Jordan curve in the diagram divides the plane into equal numbers of faces in its interior and exterior. We generate and draw all such diagrams where the total number of faces is no more than 32. We note that they are surprisingly similar to Venn diagrams in that all the faces are uniquely defined by the intersection of their curve interiors.

In Chapter 4, we show that Venn diagrams are a subset of these *balanced* diagrams. Using an exhaustive computer search, we reproduce some of the results in [18] and [10], correcting the total number of convex Venn diagrams with five curves from 17 to 18. Every 5-curve Venn diagram, up to isomorphism on the sphere, is illustrated and classified. The classification includes whether the diagram can be drawn on the plane with five convex curves. We further classify each drawing by the number of faces that are bounded by i curve segments where i is between three and the number of curves. We also show a constructive heuristic that extends an irreducible simple Venn diagram with n curves to another simple Venn diagram with $n + 1$ curves. Each of the known simple Venn diagrams up to five curves are easily extendible, using this heuristic.

Chapter 5 defines a set system that helps define Venn diagrams in terms of the distribution of the faces formed by the curves, where each face is a single element of a larger set. We define a general *half-set system* as a collection of subsets of a set of 2^n elements where the intersection of k unique subsets results in a set of size 2^{n-k}. In Chapter 6, we answer some previously open problems about Venn diagrams where the curves are represented as polyominoes. Using the half set-system, we prove that such diagrams exist where each region is a single unit square. Notably we prove the existence of such diagrams confined within bounding rectangles and show two expansion constructions: one where a 4×2^c minimum area Venn diagram is expanded to a $4 \times 2^{c+3}$ diagram, the other where a $2^r \times 2^c$ minimum area Venn diagram is expanded to a $2^{r+r'} \times 2^{c+c'}$ diagram.

1.1 Basic graph definitions and Jordan curves

We follow Douglas West's [24] basic graph definitions. A *graph* G is a triple consisting of a *vertex set* $V(G)$, and *edge set* $E(G)$, and a relation that associates with each edge two vertices, not necessarily distinct, called its *endpoints* [p. 2]. Two graphs G

and H are isomorphic if there is a bijection $f\colon V(G) \to V(H)$, such that u and v are adjacent in G if and only if $f(u)$ and $f(v)$ are adjacent in H. A *directed graph* has *directed edges*, where the endpoints are an ordered pair: the first is the *tail* vertex and the second is the *head* vertex [p. 53]. A directed edge is said to be directed *from* its tail *to* its head.

An x, y *walk* [p. 20] on a graph G is a sequence $x = v_0, e_1, v_1, e_2, \ldots, e_k, v_k = y$, of vertices and edges such that, for $1 \leq i \leq k$, the edge e_i has endpoints v_{i-1} and v_i. When a graph has no more than one edge between any two vertices, a walk is completely specified by its ordered list of vertices. When $x = y$ on a walk, it is a *closed walk*; if no vertex is repeated, except the endpoints, it is a *cycle*. The *length* of a walk or a cycle is the number of its edges. A *Hamilton cycle* on a graph G is a cycle involving all vertices of G. We say G is *Hamiltonian* if there exists such a cycle on G.

A *drawing* [p. 234] of a graph G is a function f defined on $V(G) \cup E(G)$ that assigns each vertex v a point $f(v)$ in the plane and assigns each edge with endpoints u, v a polygonal $f(u)$, $f(v)$-curve. The images of vertices are distinct. A point in $f(e) \cap f(e')$ that is not a common endpoint is a *crossing*.

1.1.1 Planar graphs

A *planar graph* [24, p. 235] is a graph that has a drawing without crossings. Such a drawing is a *planar embedding* of G. A *plane graph* is a particular embedding of a planar graph. The *faces* of a plane graph [p. 235] are the maximal regions of the plane that contain no point used in the embedding. Every face is bounded by a set of edges in $E(G)$. These edges, with their adjacent vertices, form a cycle on G. The *length* of a face in a plane graph is the length of the cycle bounding the face, which is exactly the number of edges on its boundary.

The *dual graph* [p. 236] G^* of a plane graph G is a plane graph whose vertices correspond to the faces of G. The edges of G^* correspond to the edges of G as follows: if e is an edge of G with face X on one side and face Y on the other side, then the endpoints of the dual edge $e^* \in E(G^*)$ are the vertices x, y of G^* that represent the faces X, Y of G.

Two plane graphs are *topologically equivalent* when one can be turned into the other by a *continuous transformation of the plane*. This transformation is achieved when we stretch or shrink all or parts of the plane, without tearing, twisting or

pasting it to itself. Henceforth, we consider plane graphs to be identical when they are topologically equivalent.

Any plane graph can be embedded on the sphere by a *stereographic projection*. The surface of the sphere touches the plane at its south pole S, with its north pole N diametrically opposite S. A line segment from any point x on the plane to N must intersect the surface of the sphere at a unique point x'. Figure 1.1 illustrates the relationship between point x on the plane and point x' on the sphere. Conversely, any graph that is embedded on the sphere can be embedded on the plane. Choose N so that it does not lie on a vertex or edge of the graph and follow the line from N through x' on the surface of the sphere to its endpoint x on the plane.

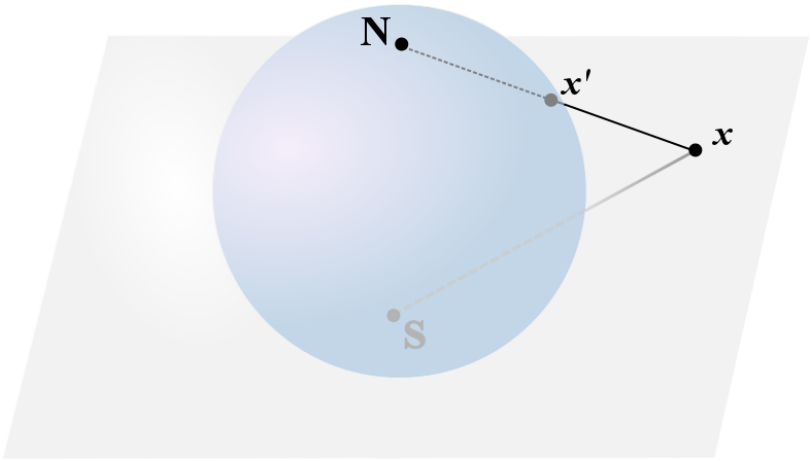

Figure 1.1: Stereographic projection of a point x on the plane and x' on the sphere.

The outer face of the plane graph that is mapped from a spherical embedding is dependent on the choice for the position of N on the sphere. A graph G embedded on the sphere is a graph, together with a set $L = \{L(v)\}$ of ordered circular lists of edges incident to each vertex $v \in V(G)$. If G is a graph embedded on the sphere, then its *mirror image* $M(G)$ is obtained by reversing all of the lists in L, alternatively described as turning the sphere inside out. We say that two plane graphs G_1, G_2 are *isomorphic on the sphere* if they are identical to a spherical embedding or its mirror image.

1.1.2 Jordan curves

A *simple closed curve* in the plane is a non-self-intersecting curve, which, by a continuous transformation of the plane is identical to a circle. It is often called a *Jordan*

curve, so named for the Jordan theorem that states that every simple closed planar curve separates the plane into a bounded interior region and an unbounded exterior [17]. In this thesis, we are interested in Jordan curves where each intersection point is the transversal crossing of exactly two curves. We refer to such a drawing of intersecting Jordan curves as *simple*. Clearly, each pair of curves intersect each other at an even number of points.

A drawing of a set of simple intersecting Jordan curves is a plane graph, whose vertices are the intersections of the curves and whose edges are the line segments connecting these vertices. The dual graph of a set of simple intersecting Jordan curves is a *quadrangulation*, a graph where all faces have length four. The terms "drawing" and "graph" are used interchangeably; the meaning will be clear from the context.

A *region* of a Jordan curve drawing with n labelled curves is identified by the unique set $X_1 \cap X_2 \cap \cdots \cap X_n$, where X_i is either the bounded interior or the unbounded exterior of curve C_i. Note that the union of several faces in the plane graph can make up a single region. When each of the $X_1 \cap X_2 \cap \cdots X_n$ sets is non-empty, the drawing is called an *independent family* of curves [15]. When each set is either empty or identified by a single face, the drawing is called an *Euler diagram*.

1.2 The Venn diagram

We follow Grünbaum's definition of a *Venn diagram* as a collection of Jordan curves $C = C_1, C_2, \ldots, C_n$ drawn on the plane such that each two curves intersect at a finite number of points and each of the 2^n sets $X_1 \cap X_2 \cap \cdots \cap X_n$ is a nonempty and connected region where X_i is either the bounded interior or unbounded exterior of C_i [16]. In short, a Venn diagram is both an Euler diagram and an independent family. As with Jordan curve drawings, a *simple* Venn diagram curve intersects only one other curve at a point, while a non-simple Venn diagram intersection point may involve more than two curves. In this thesis, we are mostly concerned with simple Venn diagrams and the reader may assume that "simple" is inferred when omitted.

Two Venn diagrams are *isomorphic* if, by continuous transformation of the plane, one of them can be changed into the other or its mirror image [20]. Note that Venn isomorphism differs from graph isomorphism. When two Venn diagrams are isomorphic on the sphere, they are said to belong to the same *class* [9]. There is only one class of simple Venn diagrams with n curves, for $n = 1, 2, 3, 4$. In Chapter 4, we

verify that there are 20 classes for $n = 5$.

Unlike a general drawing of intersecting Jordan curves, the regions of a Venn diagram are connected, so the term "face" is synonymous with "region". Each region has associated with it a unique subset of $1, 2, \ldots, n$; if the region lies interior to curve C_i, then i is in its subset. Each region also has associated with it, a *weight*, which is the cardinality of the representative subset. Note that in a Venn diagram, each edge borders exactly two regions whose weights differ by exactly one. In this thesis, we use "region" when referring to a face that is uniquely identified by its associated subset and "face" when we are not concerned with its subset.

A *family of intersecting closed curves* [4], or FISC, is the set of n intersecting Jordan curves with the property that there is a region whose weight is equal to n. Every Venn diagram is a FISC.

Monotone Venn diagrams

A Venn diagram is *monotone* if every region with weight $0 < k < n$ is adjacent to a region with weight $k - 1$ and a region with weight $k + 1$ [6]. From this definition, every region of weight one must be adjacent to the exterior region, meaning that every curve must have a segment on the boundary of the outer region. A Venn diagram is said to be *exposed* if each of its curves has a segment on the boundary of the outer region. A non-exposed Venn diagram with n curves has $k < n$ curve segments on the boundary of the outer region.

Because all the faces with weight equal to one are adjacent to the outer face whose weight is zero, a monotone Venn diagram is always an exposed diagram. Since a similar requirement holds for the common interior region with weight n, a monotone Venn diagram always has a *twin* diagram in the same class that is a mapping from the sphere with the poles reversed. If they are isomorphic, then they are called *polar symmetric*.

A Jordan curve is *convex* if any two interior points can be joined by an interior line segment. A Venn diagram is a convex diagram if all its curves are convex. Every monotone Venn diagram is isomorphic to a convex diagram [4]. Thus, in this thesis, we use monotonicity to justify that a Venn diagram is convex. A monotone Venn diagram is only isomorphic to another monotone Venn diagram. A *potentially monotone* Venn diagram belongs to the same class as a monotone Venn diagram.

1.2.1 Reducing and extending the Venn diagram

A Venn diagram V with n curves can be *extended* if a "suitable" curve C^* can be added to obtain a Venn diagram with $n+1$ curves. In order to account for all the unique regions, C^* must *split* every region of V into two pieces [25], one piece will become a region in the interior of C^*, the other a region of the exterior. Clearly, the new curve C^* is equivalent to a Hamilton cycle on the dual graph $D(V)$. The resulting Venn diagram with the addition of C^* is *reducible* because the removal of C^* clearly results in the original Venn diagram with n curves. In fact, any curve that has 2^{n-1} intersection points on a Venn diagram with n curves can be the candidate for removal on a reducible Venn diagram. If there is no such curve, then the diagram is *irreducible*.

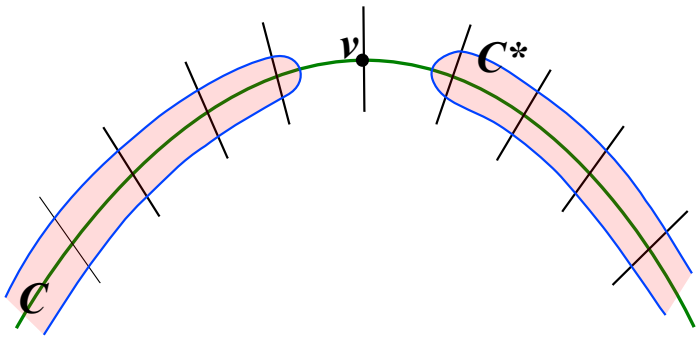

Figure 1.2: Closeup of Venn's construction

In [23], John Venn outlined a construction for Venn diagrams with any number of curves, by successively adding curves. Peter Winkler [25] proved that Venn's construction can be applied to any reducible Venn diagram by demonstrating that the new curve is a Hamilton cycle on the dual graph. On the Venn diagram, the new curve C^* follows both sides of a reducible curve C up to but not including a single arbitrary vertex v. On either side of v, C^* crosses C to connect the inside and outside C portions of C^*. Figure 1.2 shows a close up of a section where a new blue curve C^* follows the green curve C, crossing C on either side of v. The black line segments crossing C represent some other curves on the diagram. The internal region of C^* is shaded light red, and the combination of C^* and its interior looks like a thick highlighter mark.

If none of the n curves on a simple Venn diagram V has 2^{n-1} intersection points, it is not generally known whether the diagram is extendible by an additional curve. Winkler's 29 year old conjecture [25] states that it is.

Chapter 2

Jordan Curves

In this chapter, we state some facts about Jordan curve drawings as a basis for the more restrictive Jordan curve drawings discussed in Chapters 3 and 4. When John Venn defined his symbolic logic, he illustrated the concepts using Jordan curves [23]. In this thesis, we are concerned with Jordan curves where the underlying plane graph is connected and 4-regular. In other words, the Jordan curves are simple.

Proposition 2.0.1. *Let J be a Jordan curve drawing with f faces and $n > 1$ curves. Then the removal of a curve C in J that has x intersections results in a diagram with $f - x$ faces.*

Proof. Consider the drawing J, with C removed. Let the number of faces in $J \setminus C$ be j. Tracing C back onto this drawing, it is easy to see that for every intersection point as C crosses a curve in $J \setminus C$, it enters and exits an existing face, splitting that face into two faces. Therefore every face split by C is associated with two intersection points on C. Since every intersection point on C is both an entrance to one face and an exit for another, the number of intersection points of C equals the number of faces that are split. Hence x new faces are added to $J \setminus C$, so $x + j = f$ and $j = f - x$. □

Given the underlying graph of a Jordan curve drawing J with v vertices (intersections), e edges (curve segments) and f faces, the following equations, (2.0.1) and (2.0.2) are true: Firstly,

$$e = 2v, \tag{2.0.1}$$

because each vertex in J has degree four. Substituting Equation (2.0.1) into Euler's

relation, we get $f + v = 2v + 2$. Thus secondly,

$$v = f - 2. \tag{2.0.2}$$

2.1 Enumerations of simple connected collections of Jordan curves

Recall that the dual graph $D(J)$ of a Jordan curve drawing J is a quadrangulation. We can traverse each curve of J on $D(J)$ by traversing the *cross edges*: the dual edges that cross the curve C in J. These cross edges are never adjacent on the boundary of each face in $D(J)$. See Figure 2.1 for an example of a dual graph of a Jordan curve drawing. In contrast, Figure 2.2 shows a dual graph where the some cross edges are adjacent on a face, and thus the primal graph has a self-intersecting curve.

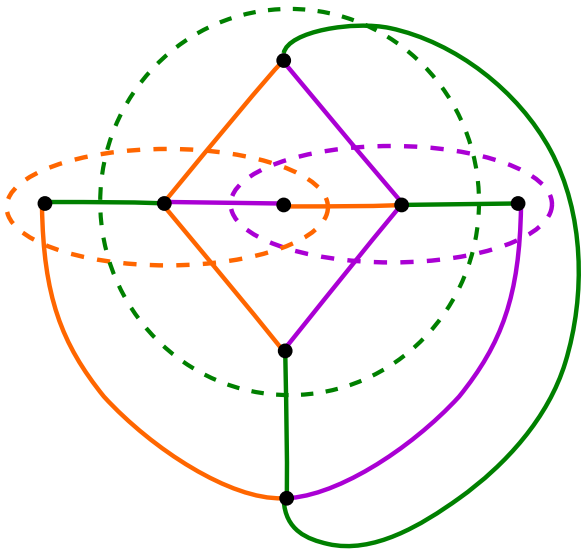

Figure 2.1: A dual quadrangulation (solid lines) where the cross edges are coloured to match the corresponding Jordan curve in the primal graph (dashed lines). No edges of the same colour are adjacent on the boundary of a dual face.

We used Brinkman and Mackay's planar graph generating program, `plantri` [3], described in [2] as an "isomorph-free" generator of many classes of planar graphs. We set it to generate quadrangulations that are 2-connected for graphs with n vertices for $4 \leq n \leq 24$. Each graph produced by `plantri` was further filtered through a plugin that checks for valid cross edges. See the algorithm in Figure 2.3, which outlines this plugin. If a quadrangulation is accepted by the plugin, the resulting

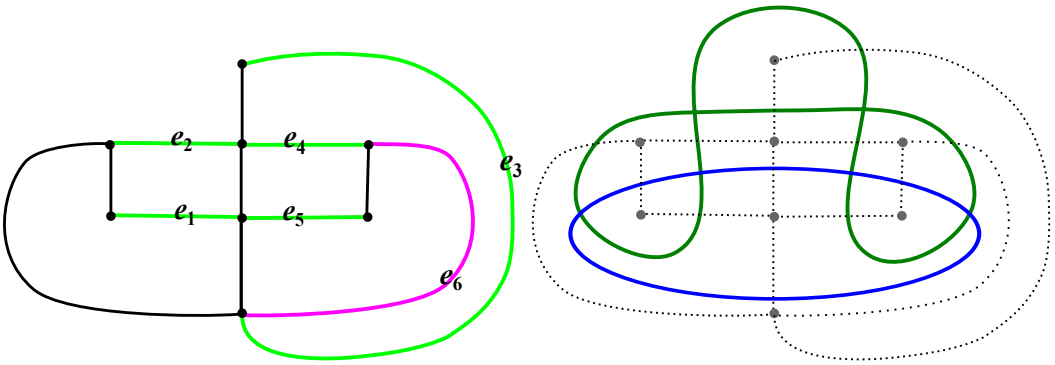

Figure 2.2: A rejected quadrangulation. The green edges e_1 to e_5 show cross edges. However, edge e_6 is adjacent to e_3 and e_4 and cannot form a matching in either of its adjacent faces. The illustration on the right shows the dual of the quadrangulation; the green curve is self-intersecting.

graph is the dual graph of a single Jordan curve drawing. Here after, in this section and in Chapter 3, we use the simpler term *diagram* with the understanding that two diagrams are isomorphic if they are isomorphic on the sphere.

In plantri, each edge e consists of two directed edges $\vec{e_1}$ and $\vec{e_2}$; the inverse of $\vec{e_1}$ is $\vec{e_2}$ and vice versa. Each directed edge $\vec{e_i}$ has a previous and next directed edge, $\text{prev}(\vec{e_i})$ and $\text{next}(\vec{e_i})$, determined by the cyclic ordering of edges around the tail vertex.

All accepted graphs are counted and the numbers shown in Table 2.1. The columns classify the diagrams by the number of curves, while the rows classify them by the number of faces.

Input:

- graph G: a 2-connected quadrangulation with n vertices, stored as an array of unmarked edges.

Output:

- accept if G is the dual of a Jordan curve drawing, reject if not.
- the number of curves.

```
1  nextEdge ← an unmarked edge of G;
2  i ← 1;
3  curveCount ← 0;
4  while nextEdge is not null do
5      e ← nextEdge;
6      repeat
7          mark e as part of curve i;
8          if prev(e⃗1), next(e⃗1), prev(e⃗2) or next(e⃗2) are marked as part of curve i
           then
9              return false ;      // the Jordan curve is self-intersecting
10         end
11         e ← new edge opposite e on the 4-face;
12     until e = nextEdge ;
13     curveCount ← curveCount +1;
14     nextEdge ← an unmarked edge of G;
15     i ← i + 1;
16 end
17 return true, curveCount
```

Figure 2.3: Algorithm: Determine whether a quadrangulation is the dual of a simple, connected Jordan curve drawing.

Table 2.1: The numbers of non-isomorphic Jordan curve diagrams with f faces and n curves.

$f \setminus n$	2	3	4	5	6	7	8	9	10	11
4	1									
6	1	1								
8	2	2	1							
10	4	5	1	1						
12	13	26	9	1	1					
14	45	181	98	11	1	1				
16	212	1462	1245	220	14	1	1			
18	1165	13,990	17,441	4857	418	17	1	1		
20	7649	145,034	251,397	104,734	13,767	709	21	1	1	
22	55,423	1,593,666	3,633,555	2,120,412	418,745	32,212	1131	24	1	
24	435,913	18,215,569	52,555,710	40,554,089	11,340,049	1,301,628	66,969	1692	28	1

Chapter 3

Face-balanced Curves

Many facts about Venn diagrams have been determined by examining the underlying graph and the dual of the Venn diagram. Winkler's conjecture [25] is equivalent to stating that the dual of a Venn diagram always contains a Hamilton cycle. Chilakamarri and Hamburger [8] proved, using graph theory, that every Venn diagram of n curves can be extended to a non-simple Venn diagram of $n+1$ curves by the addition of a suitable Jordan curve.

We note the following three necessary properties of a Venn diagram as a graph:

- The Venn diagram as a plane graph is an embedding of a 4-regular planar graph.
- Every curve divides the plane, with 2^{n-1} regions interior and 2^{n-1} exterior.
- Every face on the Venn diagram has unique curve segments on its boundary.

In this chapter, we examine the set of diagrams that share these properties with the Venn diagrams. We also classify all such diagrams with up to 32 faces.

3.1 Introductory results

Definition 3.1.1. A drawing of Jordan curves on the plane is *balanced* if it has the property that, for each curve, the number of interior faces and exterior faces are equal.

Proposition 3.1.2. *Any two Jordan curves in a balanced drawing must intersect.*

Proof. Consider a balanced drawing with $2k$ faces and a curve C_i that does not intersect curve C_j. Clearly if the interior of C_i is entirely contained in the interior of

C_j, then the number of faces interior to C_j includes all of the faces interior to C_i, plus the one extra that is exterior to C_i. If the intersection of their interiors is the empty set, then the number of faces exterior to C_i must include both the interior faces of C_j plus the exterior unbounded face. In both cases, the number of interior faces are not equal and the drawing is not balanced. □

Proposition 3.1.3. *Every drawing with two intersecting Jordan curves is a balanced drawing.*

Proof. By Equation (2.0.2), two Jordan curves that intersect each other $2x$ times will divide the plane into $2x+2$ faces. Tracing along one of the curves, C_1, one can easily see that the intersections alternate from the interior to the exterior of C_2. So C_2 has $x+1$ interior faces and $x+1$ exterior faces. □

Suppose we have a balanced Jordan curve drawing D, with $n > 1$ curves. For any curve C_i on D, let $|C_i|$ be the number of faces in the interior of C_i and $|\overline{C_i}|$ be the number of faces on the exterior of C_i.

Lemma 3.1.4. *Let C_i and C_j be two distinct Jordan curves in a balanced drawing D; then*

$$|C_i \cap C_j| = |\overline{C_i} \cap \overline{C_j}|$$

Proof. Let $2k$ be the total number of faces in D. Then $|C_i| = |C_j| = |\overline{C_i}| = |\overline{C_j}| = k$. By basic set theory:

$$\begin{aligned} |C_i \cap C_j| &= |C_i| + |C_j| - |C_i \cup C_j| \\ &= 2k - |C_i \cup C_j| = |\overline{C_i \cup C_j}| = |\overline{C_i} \cap \overline{C_j}| \end{aligned}$$

□

At this time, we cannot say that, for any balanced Jordan curve drawing, the intersection of k curves, where $2 < k \leq n$, must be a non-empty region. We can say that every known balanced Jordan curve drawing, including the sets catalogued in Section 3.2, has a nonempty region for the intersection of all n of the curve interiors. We make the following conjecture.

Conjecture 3.1.5. *Every balanced Jordan curve drawing is a family of intersecting closed curves, FISC [4].*

We are interested in a restricted version of the balanced Jordan curve drawing that is more closely related to the Venn diagram. Consider a drawing where none of the faces is bordered by more than one segment of the same curve.

Definition 3.1.6. A balanced drawing is called a *face-balanced Jordan curve drawing* or *face-balanced drawing* if each of the faces is bordered by no more than one segment of the same curve.

Recall that any drawing of two intersecting Jordan curves is a balanced diagram. However, the only diagram that is *face-balanced* is the Venn diagram with two curves, since each face consists of a single segment of each of the curves.

Proposition 3.1.7. *A face-balanced drawing consisting of $n \geq 3$ curves must have at least three curve segments bordering each face.*

Proof. Suppose a face-balanced drawing exists, with face f bounded by two segments s_1 and s_2, from curves C_1 and C_2, respectively. Consider the neighbouring face, f', that shares segment s_1 with f. The curve C_2 bounds f' at both endpoints of s_1. If no other curve is on the boundary of f', then C_2 has two interiors, and the drawing is an embedding of the Venn diagram with two curves. By Proposition 3.1.2, when $n \geq 3$, there must be at least one other curve C_3 that intersects C_2. Since it cannot border f, C_3 borders f', and must split C_2 into two segments on f', as illustrated in Figure 3.1. However, the existence of two segments of C_2 in f' contradicts Definition 3.1.6. □

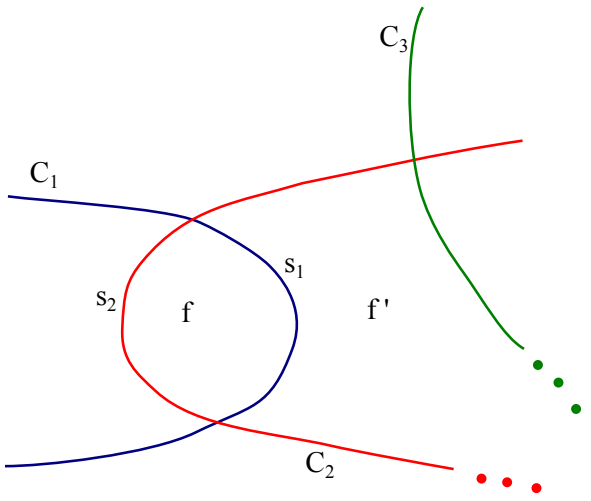

Figure 3.1: There can be no border of two edges on a face-balanced diagram.

Note that a curve in a face-balanced drawing with n curves can have no more than n curve segments bordering any face. We know something about the number of pairwise intersections of curves on a face-balanced drawing.

Lemma 3.1.8. *Let x be the number of intersections of distinct curves C_i and C_j on a face-balanced drawing of three or more curves. Then both C_i and C_j each have at least $2x$ intersection points.*

Proof. Suppose that distinct curves, C_i and C_j, of a face-balanced drawing intersect each other x times. Suppose also, that on a walk around curve C_i we encounter three consecutive intersections of curve C_j. Without loss of generality, consider the face f, on the right as we walk along C_i, between the two intersections of C_j: By Proposition 3.1.7, this face must be bounded by at least one other curve C_k that intersects C_j on the border of f. Thus C_j has two segments on the border of f, and by Definition 3.1.1, the drawing is not face-balanced. By the supposition, there cannot be consecutive intersections of C_j on C_i. Therefore, each of C_i and C_j must have a total number of intersection points that is at least $2x$. □

Lemma 3.1.9. *In a face-balanced drawing D with n curves and $2k$ faces, a curve C_i has no more than k total intersection points.*

Proof. Suppose C_i has $x > k$ intersections. Then by Proposition 2.0.1, the removal of C_i results in a diagram $D \setminus C_i$ with less than k faces. Tracing C_i on D over its $x > k$ intersection points, we must enter and exit a face in $D \setminus C_i$ more than once. However, this means that there is a face in D that contains more than one curve segment of C_i. By Definition 3.1.6, D cannot be a face-balanced drawing when C_i has more than k intersection points. □

Corollary 3.1.10. *A curve C of a face-balanced drawing of n Jordan curves and $2k$ faces has x total intersection points, where $2(n-1) \leq x \leq k$.*

Proof. The lower bound for x is obtained from Proposition 3.1.2 and the upper bound is obtained from Lemma 3.1.9. □

Corollary 3.1.11. *Let x_{ij} be the number of intersections of curves C_i and C_j on a face-balanced drawing on $n \geq 3$ curves with $2k$ faces. Then x_{ij} is an even number and*

$$2 \leq x_{ij} \leq \frac{k}{2}.$$

Proof. Clearly, the number of intersections of two Jordan curves is an even number. Proposition 3.1.2 gives us $x_{ij} \geq 2$. By Lemma 3.1.9, a curve has no more than k total intersection points, and by Lemma 3.1.8, the number of intersections of two curves must be one half that value, so $x_{ij} \leq \frac{1}{2}k$. □

3.2 Face-balanced drawings with up to 32 faces

We now find all face-balanced drawings with 32 or fewer faces. Using the computer search described in Section 2.1, we added two more requirements to the plugin outlined in Algorithm 2.3. These are as follows:

1. Every set of cross edges must form a matching on the dual graph.
2. Each set of cross edges, if deleted from the dual graph, would leave a disconnected graph with two components; each of these components having an equal number of vertices.

The first requirement guarantees that no face contains more than one Jordan curve segment on its boundary. This differs from a general Jordan curve drawing where the dual edges cannot be adjacent on a face boundary, but still adjacent on the larger graph. For example, Figure 2.1 has two adjacent orange edges and two adjacent purple edges, which is disallowed for face-balanced drawings. Chilakamarri et al. [10] state that the sets of matchings are clearly present on the dual of every Venn diagram. It is just as clear that they are present on every face-balanced drawing. Recall that the vertices on each of the cross edges represent the faces of a Jordan curve drawing. If two cross edges that correspond to the same curve on the drawing meet at a single vertex in the dual graph, then two edges of the same curve are present on a face of the drawing.

The second requirement is a clear requirement of the face-balanced drawing.

These two added requirements validate Venn diagrams as well as face-balanced curves. However, we leave the set of Venn diagrams until Chapter 4. The non-Venn diagrams resulting from the computer search, illustrated in Figures 3.2 to 3.10 have interesting properties: Every diagram has an exposed embedding on the plane. In each plane embedding of a diagram, each of the faces is uniquely defined by some $X_1 \cap X_2 \cap \cdots \cap X_n$, where X_i is either the bounded interior or unbounded exterior of curve C_i. Several of the monotone diagrams have an *antipodal symmetry*, antipodal

referring to two points on a sphere whose distance is as great as possible. On a Jordan curve drawing, the antipodal points we are concerned with are those where the same two curves intersect. The *symmetry* occurs when every point has an antipodal mate. An easy way to check for this in a diagram is to walk around a curve C_i, and name each curve that intersects C_i. If the sequence can be divided into two repeating subsequences, then all of C_i's crossings are antipodal.

A collection of simple closed Jordan curves in the plane is called a family of *pseudo-circles* if any two curves intersect transversally either twice or zero times [1]. Several of the diagrams are families of pseudo-circles, easily identifiable because each of the n curves has $2(n-1)$ intersections, the lower limit, by Corollary 3.1.10. Figure 3.9 is the only drawing that does not have antipodal symmetry. The pseudo-circle drawings that have antipodal symmetry can all be drawn on the sphere as great circles.

Table 3.1 summarizes information on all the face-balanced diagrams that are not Venn diagrams. Each diagram is categorized by the following information:

- The figure number and link of the illustrated drawing in this section.
- The number of faces.
- The number of Jordan curves in the diagram.
- The weight sequence, $w_0, w_1, \ldots, w_n$ where w_k is the number of faces that have weight k. Note that the weight sequence can vary, depending on the drawing. Of note is the fact that the weight sequence of a Venn diagram is a list of the binomial coefficients. Some of the face-balanced drawings have palindromic weight sequences. Some have increasing, then decreasing sequences, some do not. All have non-zero values.
- Whether the diagram has a monotone embedding [m], has antipodal symmetry [a] and is a family of pseudo-circles [p]. A 'y' value indicates *yes* and a 'n' value indicates *no*.
- The listing of the faces as a tuple $k_3, k_4, \ldots, k_n$, where n is the number of curves and k_i is the number of faces bordered by i curve segments.

Based on the observations of this group of face-balanced diagrams, we make the following conjecture:

Table 3.1: All the face-balanced diagrams with 32 faces or less. Three diagrams have 28 faces and 5 curves, five have 32 faces and 6 curves.

figure	curves	faces	weight sequence	m/a/p	face lengths
3.2	4	14	$1, 4, 4, 4, 1$	y/y/y	$8, 6$
3.3	5	22	$1, 5, 5, 5, 5, 1$	y/y/y	$10, 10, 2$
3.4	5	26	$1, 5, 7, 7, 5, 1$	y/y/n	$14, 6, 6$
3.5	5	28	$1, 5, 8, 8, 5, 1$	n/n/n	$14, 8, 16$
3.6	5	28	$1, 5, 7, 10, 4, 1$	y/n/n	$12, 12, 4$
3.7	5	28	$1, 5, 8, 8, 5, 1$	y/n/n	$12, 12, 4, 0$
3.8	6	32	$1, 5, 5, 10, 5, 5, 1$	y/y/y	$3, 0, 12, 0$
3.9	6	32	$1, 4, 7, 8, 7, 4, 1$	y/n/y	$3, 24, 0, 0$
3.10	6	32	$1, 5, 6, 8, 6, 5, 1$	y/y/y	$14, 12, 6, 0$
3.11	6	32	$1, 6, 6, 6, 6, 6, 1$	y/y/y	$12, 18, 0, 2$
3.12	6	32	$1, 5, 7, 6, 7, 5, 1$	y/y/y	$12, 16, 4, 0$

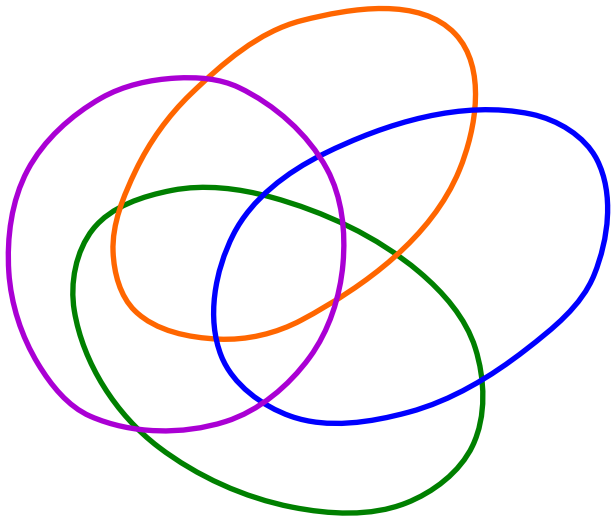

Figure 3.2: Face-balanced, 4 curves, 14 faces, monotone, antipodal symmetry, pseudo-circles.

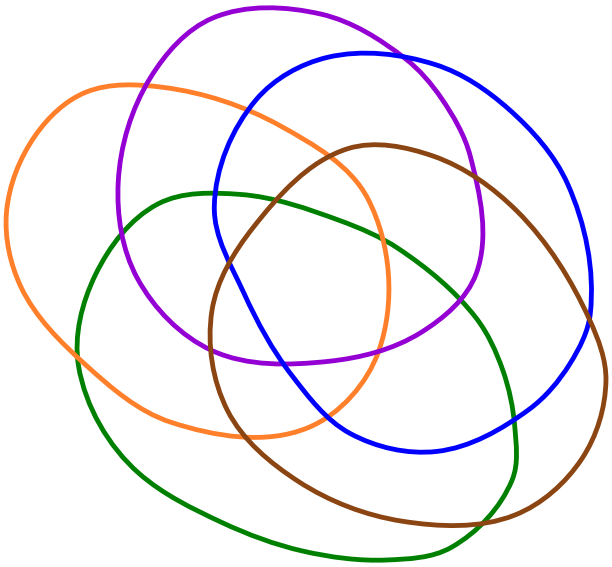

Figure 3.3: Face-balanced, 5 curves, 22 faces, monotone, antipodal symmetry, pseudo-circles.

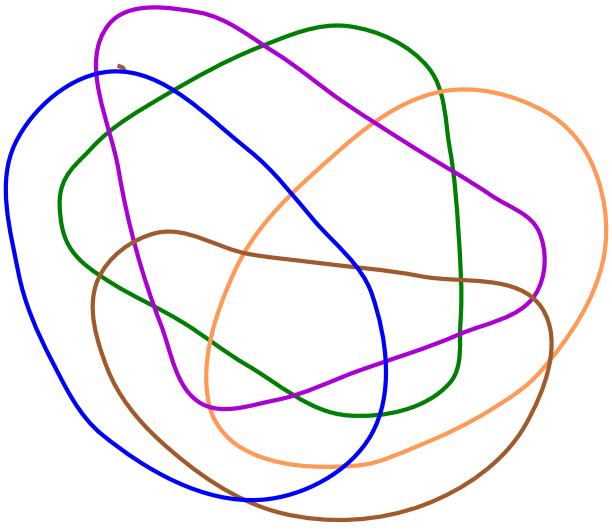

Figure 3.4: Face-balanced, 5 curves, 26 faces, monotone, antipodal symmetry.

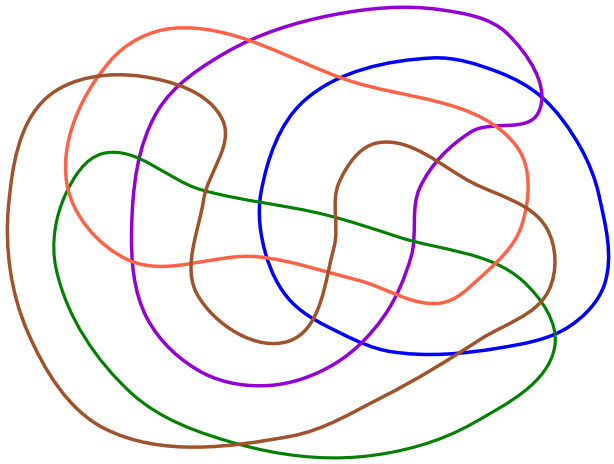

Figure 3.5: Face-balanced, 5 curves, 28 faces.

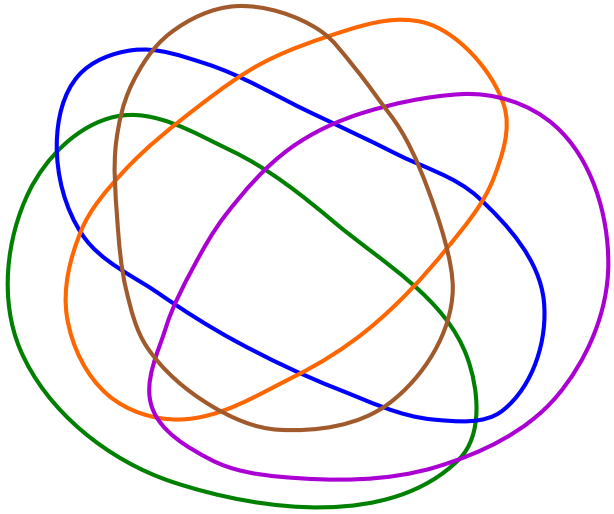

Figure 3.6: Face-balanced, 5 curves, 28 faces, monotone. Intersection points on the orange, brown and blue curves have antipodal symmetry.

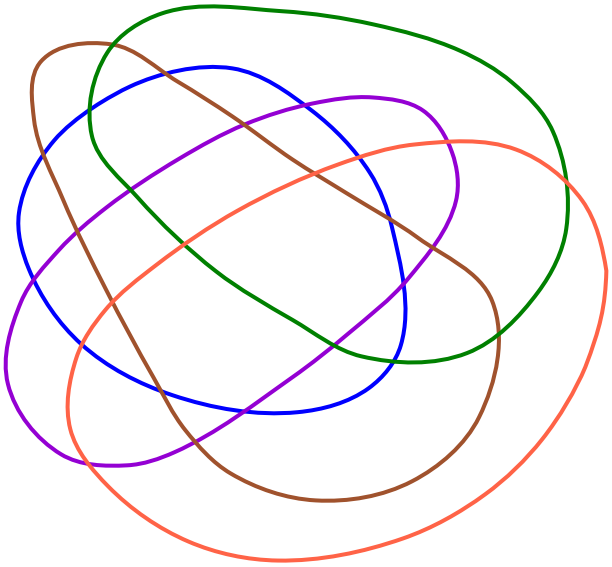

Figure 3.7: Face-balanced, 5 curves, 28 faces, monotone. In this drawing each curve sequence is made up of 2 mirrored sequences.

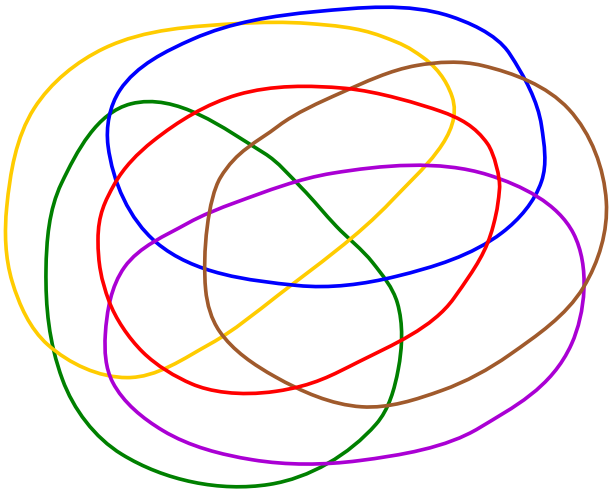

Figure 3.8: Face-balanced, 6 curves, 32 faces, monotone, antipodal symmetry, pseudocircles.

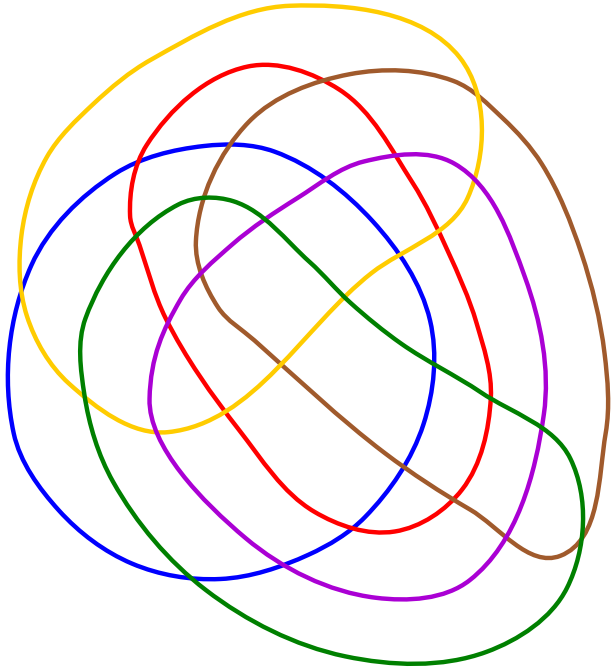

Figure 3.9: Face-balanced, 6 curves, 32 faces, monotone, pseudo-circles.

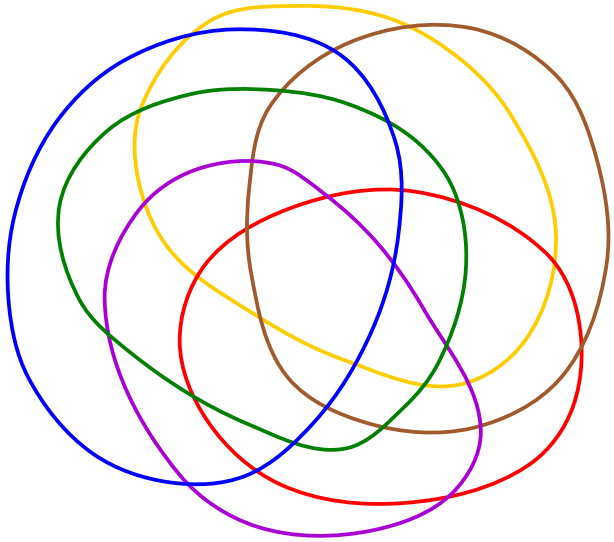

Figure 3.10: Face-balanced, 6 curves, 32 faces, monotone, antipodal symmetry, pseudo-circles.

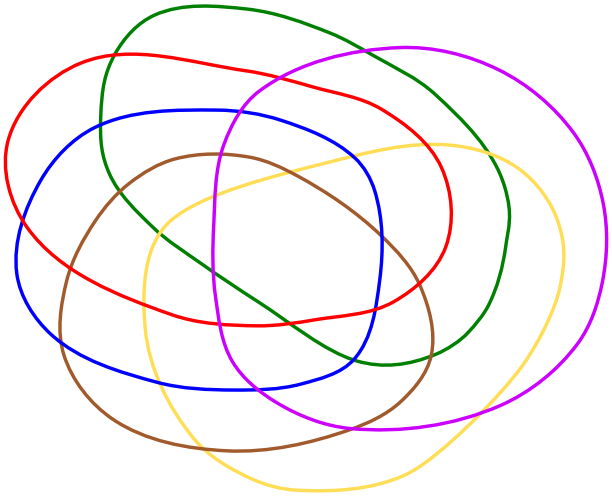

Figure 3.11: Face-balanced, 6 curves, 32 faces, monotone, antipodal symmetry, pseudo-circles.

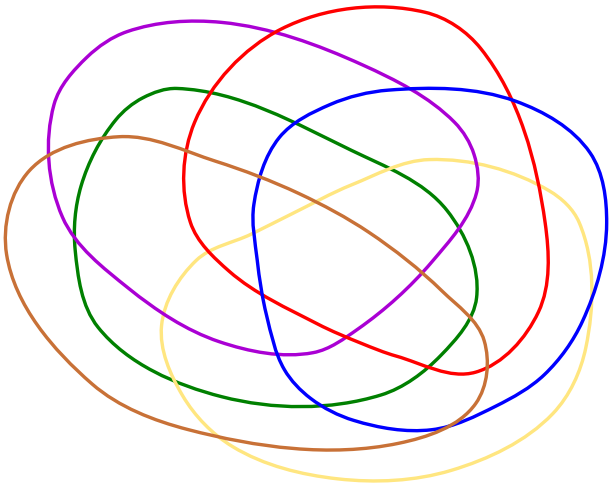

Figure 3.12: Face-balanced, 6 curves, 32 faces, monotone, antipodal symmetry, pseudo-circles.

Conjecture 3.2.1. *In a face-balanced drawing, each set $X_1 \cap X_2 \cap \cdots \cap X_n$ is either empty or is a connected region where X_i is either the bounded interior or the unbounded exterior of curve C_i.*

If Conjecture 3.2.1 is true, then face-balanced diagrams are a subset of Euler diagrams.

To help prove or disprove Conjecture 3.2.1, we may need to know how the number of regions and the number of curves are related. If it is possible to create a face-balanced drawing with n curves and greater than 2^n regions, then Conjecture 3.2.1 is false. We make the following conjecture of a lesser result.

Conjecture 3.2.2. *In a face-balanced drawing with n curves, the number of regions does not exceed 2^n.*

The facts that we currently know about face-balanced diagrams, beyond the set that was generated, are derived from basic planar graph theory. The face-balanced drawing of n curves and $2k$ faces is a 4-regular graph, where each k-face satisfies $3 \leq k \leq n$. The following lemma considers such a graph.

Lemma 3.2.3. *Consider a 4-regular planar graph G, with $2k$ faces, each of whose length is between 3 and some $n \geq 3$. At least eight faces have length three and the average face length in G is $4 - \frac{4}{k}$.*

Proof. Let f_i be the number of faces on G that have length i, noting that $3 \leq i \leq n$. Then

$$\sum_{i=3}^{n} f_i = 2k. \tag{3.2.1}$$

Because each edge borders exactly two faces, and by Equations (2.0.1) and (2.0.2)

$$\sum_{i=3}^{n} i f_i = 2e = 4v = 4(2k-2) = 8k-8. \tag{3.2.2}$$

The average face size is $\frac{2e}{2k} = \frac{8k-8}{2k} = 4 - \frac{4}{k}$. To calculate the minimum number of

three faces:

$$\begin{aligned} 3f_3 &= 8k - 8 - \sum_{i=4}^{n} if_i, \text{ from Equation (3.2.2)} \\ &\leq 8k - 8 - 4\sum_{i=4}^{n} f_i \\ &\leq 8k - 8 - 4(2k - f_3), \text{ by Equation (3.2.1)} \\ f_3 &\geq -8k + 8 + 8k \\ &\geq 8. \end{aligned}$$

□

Corollary 3.2.4. *The only face-balanced diagram with three curves is the only face-balanced diagram with eight faces. It is isomorphic to the well-known Venn drawing with three circles.*

Proof. The computer search for all face-balanced diagrams with eight faces produced only one result, a three curve diagram.. By Corollary 3.1.11, every face on the graph representation of a 3-curve face-balanced drawing has length three. Since by Lemma 3.2.3, the average face size $= 4 - \frac{4}{k} = 3$, then $k = 4$ and the number of faces is eight. The Venn diagram on three curves is face-balanced, and has three faces. □

3.3 Reducibility and extendibility of face-balanced curves

A face-balanced drawing with n curves is *reducible* if the removal of any one of its curves results in a face-balanced drawing of $n - 1$ curves. Likewise, a face-balanced drawing with n curves is *extendible* if the addition of a suitable curve results in a face-balanced drawing of $n + 1$ curves.

If Conjecture 3.2.1 is proven to be untrue, then it is conceivable that the removal of a curve with k intersection points will result in a Jordan curve drawing where two curve segments from the same curve border a single face. However, if the conjecture is true, then we show that it is possible to reduce a face-balanced drawing by removing such a curve.

We first demonstrate a fact that is shown to be true for Venn diagrams by Chilakamarri et al. in [9]; however, we use a different approach.

Lemma 3.3.1. *If the faces on a balanced curve drawing are all unique regions on the plane, then each face has no more than one segment of a single curve as part of its boundary.*

Proof. Suppose every face on a balanced curve drawing D with n curves is a unique region. Let f be a face on D with two distinct edges, s_1 and s_2 from curve C_i, as part of its boundary. Without loss of generality, let f lie in the interior of C_i. Then the element i is part of the subset of f. Sharing s_1 with f as part of its boundary is the face f', identified as region $f \setminus \{i\}$. Likewise, sharing s_2 with f as part of its boundary is the face f^*, identified as region $f \setminus \{i\}$. By the lemma statement the faces $f' = f^*$.

Note now that f and f' share the same two boundaries of s_1 and s_2. Therefore no curve can run in parallel between these line segments because that curve would split either f or f', resulting in a disconnected region.

Consider the two non-empty sets of curve segments on the boundary of f that separate s_1 and s_2: the curves to which they belong must be disjoint, since a curve that has segments in both sets would have to connect through f or f'. However, by Proposition 3.1.2, the curves in the two sets must intersect each other. So there must be only one set of curves separating s_1 and s_2 in f, but then s_1 and s_2 are connected and we have a contradiction. □

Now we can define the requirements for a reducible face-balanced drawing where all the regions are connected.

Lemma 3.3.2. *If a face-balanced drawing of $2k$ connected regions and $n \geq 2$ curves contains a curve C that has k intersection points, then the drawing is reducible to a face-balanced drawing of k connected regions.*

Proof. Consider a curve C_i that has k intersection points in a face-balanced drawing D, where each face defines a unique region. By Definition 3.1.6, C_i can have no more than one segment on the boundary of each of the k faces, so it must be present on the boundary of every face in D. Let C_j be another curve in D. By definition 3.1.1, there are k faces in the interior of C_j. Since C_i has a segment on the boundary of each of C_j's interior faces, C_j has $k/2$ internal faces in $D \setminus C_i$. By Proposition 2.0.1, the removal of C_i results in a diagram with k faces, meaning that C_j contains half the total faces of $D \setminus C_i$ in its interior.

Since each face defines a region in D, every pair of faces that share C_i as part of their boundary differ by a single element i, where $1 \leq i \leq n$. When C_i is removed, i is no longer an element in any of the regions, so not only are the pairs merged into a single face, they become the same region. Since the faces of $D \setminus C_i$ are unique regions, by Lemma 3.3.1 every face in $D \setminus C_i$ has unique curve segments on its boundary and is face-balanced. □

The generated set of face-balanced diagrams has interesting and promising properties. In our search for face-balanced drawings among all quadrangulations with n curves and $v \leq 32$ vertices, the result is a Venn diagram whenever the number of vertices is a power of two. It will be interesting to find more of them for further study using an algorithm that generates all face-balanced curves directly.

Chapter 4

Venn Curves

We are familiar with Venn diagrams by the ubiquitous 3-circle drawing and more recently by the beautifully rendered symmetric diagram called Newroz [19]. This chapter discusses Venn diagrams in three sections, each of which can be read independently. In the first section, we demonstrate that n curve Venn diagrams are a subset of face-balanced curves. In the second, we look at certain Venn diagrams with particular face structures that make them easily extendible. In the third section, we verify the known simple Venn diagrams with five curves, using the same plugin to `plantri` [3] that tests for face-balanced diagrams in Section 3.2. This confirms the set of all known classes of Venn diagrams for $n = 1, 2, 3, 4, 5$.

4.1 Venn diagrams are face-balanced diagrams

By definition, the n Jordan curves of a Venn diagram divide the plane into 2^n unique connected regions, and each curve has exactly 2^{n-1} interior regions. We also know from Lemma 3.3.1 and from [10] that the faces on Venn diagrams do not contain more than one segment of each curve. By Definition 3.1.6, Venn diagrams are a subset of face-balanced drawings.

4.2 Extending a Venn diagram

There remains no proof or exception to Winkler's conjecture for irreducible simple Venn diagrams. On any given simple Venn diagram, we need to find the Hamilton cycle on the dual graph of the Venn diagram in order to extend it. Finding a Hamil-

ton cycle on a 3-connected bipartite planar graph is an NP-complete problem [13, p. 199]. However in this section, we describe a heuristic for finding a particular type of Hamilton cycle on a Venn diagram.

Recall from Section 1.2.1, that if a Venn diagram has a reducible curve C, then the diagram is extendible. Figure 1.2 illustrates that the new curve can be described as the boundary of a highlighter mark that covers C, stopping short on either side of an arbitrarily chosen vertex v.

4.2.1 Extending a simple irreducible Venn diagram

When none of the n curves on a Venn diagram has 2^{n-1} intersection points, we apply the following heuristic:

Step one: set up

Create a "highlighter mark" over curve C, omitting a single arbitrarily chosen vertex. Figure 4.2 below demonstrates what we mean by a highlighter mark. This exposes a set of *isolated faces*, faces that do not intersect the highlight mark. We call faces that do intersect the highlight mark *highlighted faces*. Clearly, every highlighted face is split by the new curve, the boundary of the highlight mark. Vertices are also considered highlighted if they are covered by the highlight.

Step two: try to highlight every face

We consider a vertex v as *qualifying* if

- v is incident to two adjacent highlighted faces and
- v is incident to two adjacent isolated faces.

Since the Venn graph is 4-regular, each vertex is adjacent to exactly 4 faces. Therefore, qualifying vertices are easily determinable. The heuristic is outlined in Figure 4.1. Figure 4.2 illustrates the beginning and end of step 11.

Theorem 4.2.1. *If the heuristic of Figure 4.1 ends with the set H containing every face in the Venn diagram V, then V is extendible to a simple Venn diagram with $n+1$ curves. Furthermore, the boundary of the highlight mark is the new curve, C^*.*

Input:

- Venn diagram V.

Output:

- the set H of highlighted faces.

```
1  choose a curve C and a vertex v on C;
2  highlight all vertices and edges of a walk on C, EXCEPT v;
3  insert all highlighted faces into set H;
4  insert all qualifying vertices into queue Q;
5  while Q is not empty do
6      repeat
7          v ← dequeue(Q);
8      until v is a qualifying vertex ;
9      I ← the pair of isolated faces incident to v ;
10     u ← the highlighted vertex adjacent to v;
11     highlight the edge uv and v, forming a continuous highlight mark from u;
12     insert both faces of I into H;
13     insert all qualifying vertices adjacent to v into Q;
14 end
```

Figure 4.1: The heuristic for extending a simple Venn diagram

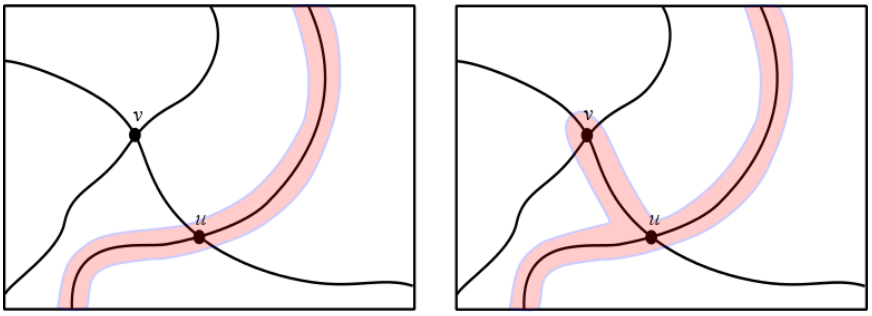

Figure 4.2: Before and after step 11 in the heuristic of Figure 4.1

Proof. The statement is true if C^* forms a Hamilton cycle on the dual $D(V)$ of the Venn diagram V. Equivalently: a) the boundary of the highlight mark is a continuous non-self-intersecting curve that b) splits every face in V exactly once.

We show that C^* is continuous because the highlight mark is continuous. We show that C^* does not self-intersect because the highlight mark contains no cycles.

We also show that a single curve segment of C^* splits every face on V.

Prior to the beginning of line 5, the highlight mark is connected, non-self-intersecting and contains no cycles. Suppose that before the k^{th} iteration of the loop beginning at line 5, the highlight is continuous, non-self-intersecting and contains no cycles. This means that every face in V is either split by a single curve segment of C^* or remains an isolated face. By the theorem statement, the heuristic continues if isolated faces remain, in which case there is a k^{th} iteration and steps $7, 9$ and 10 are implemented. Before the statement of line 11 is executed, v is not highlighted, while u is. After the statement is executed, all of u, v and the edge uv are highlighted. The faces of set I are newly highlighted. The highlight remains connected, non-self-intersecting and no cycle was introduced. Furthermore C^* has been stretched to split the faces in I while remaining connected in the faces that were already split before the k^{th} iteration of the while loop.

Because C^* meets all the necessary conditions at the end of each loop, if every face is highlighted, then C^* is a curve that extends V. □

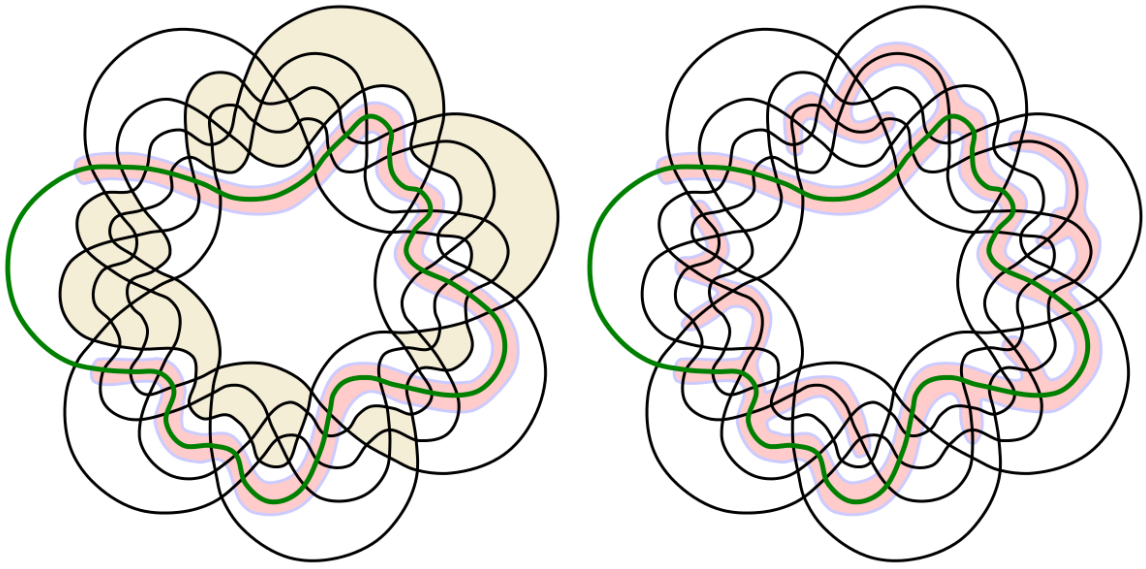

Figure 4.3: A possible extension of Ruskey's Victoria [20], using the steps in Algorithm 4.2. The shaded areas in the left drawing are the isolated faces after the highlight mark is set along the green curve. The right figure is an example of a successful completion of the algorithm. The light blue line boundary of the highlight mark is the new curve C^*.

Figure 4.3 illustrates the entire process on Victoria, Ruskey's symmetric Venn diagram from [20]. For the statement in line 1, we choose the green curve. The

left diagram shows the progress before the first loop; the isolated faces are shaded beige. Further iterations stretch the highlight mark into these beige faces until all are highlighted. The light blue exterior of the highlight is a new curve that extends Victoria.

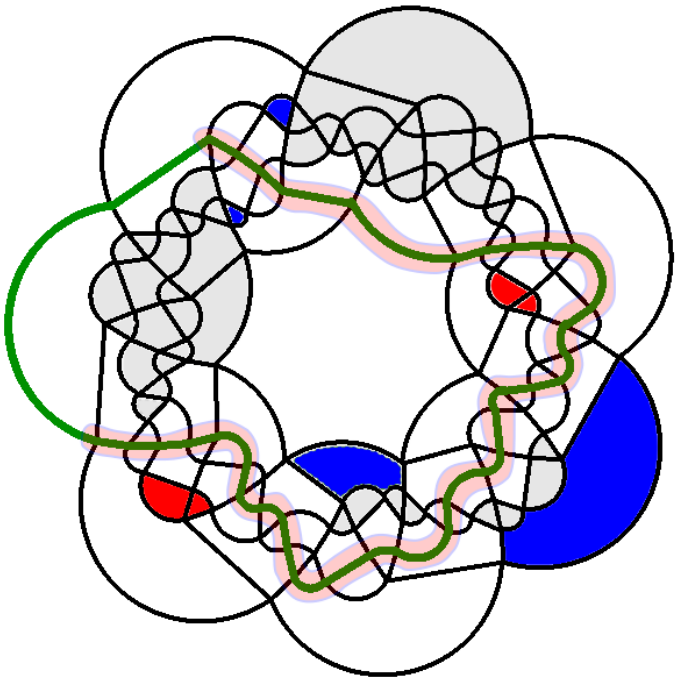

Figure 4.4: The isolated faces of Manawatu. The grey faces have qualifying vertices; the blue faces do not. The red faces cannot be reached along a crossing curve.

If, during the process, there is a face where every vertex on its boundary is not a qualifying vertex, then the heuristic will not produce an extending curve. Edward's Manawatu [12], shown in Figure 4.4, although extendible, is an example of a Venn diagram where some faces have no qualifying vertices. The faces shaded blue have no qualifying vertex on their walk because each vertex is adjacent to an odd number of isolated faces. The faces shaded red do have a common vertex that is adjacent to an even number of isolated faces; however these vertices are unreachable because they are not adjacent to a qualifying vertex.

There are, however, an infinite number of irreducible Venn diagrams that are easily extendible by the heuristic of Figure 4.1. All the irreducible Venn diagrams with 5 curves, illustrated in Section 4.3 are extended using the algorithm. In the next subsection, we show that there are infinitely many more that can be extended using this heuristic.

4.2.2 The DE property

In this subsection, we use a modification of Grünbaum's proof [15] of the existence of n-curve simple irreducible Venn diagrams for all $n \geq 5$. This constructive proof creates Venn diagrams that can all be extended using the heuristic of Figure 4.1.

Let n and k be integers such that $n > 4$ and $1 \leq k \leq n$. Let V be a simple irreducible Venn diagram with n curves with curve C_k having $2^{n-1} - 2$ intersection points. There exist exactly four faces, S_i, S_e, R_i, R_e, all of whose boundaries do not contain a segment of C_k. We say that V has a *double extendibility* property, or DE property, if the following conditions with respect to C_k, S_i, S_e, R_i, R_e and some curve C_m that is not C_k are true:

1. Faces R_i and R_e lie interior to C_k. Faces S_i and S_e lie exterior to C_k.

2. Faces R_i and R_e share a common edge that is a segment of C_m. Faces S_i and S_e also share a common edge that is a segment of C_m. Faces R_i and S_i are interior to C_m and R_e and S_e are exterior to C_m.

3. The curve segment along C_m on the boundary of both R_i and R_e has a vertex u_m that is incident to u_k, a vertex where C_m and C_k intersect on V. Likewise, the curve segment along C_m on the boundary of both S_i and S_e has a vertex v_m that is incident to v_k, another vertex where C_m and C_k intersect on V.

Figure 4.5 shows a 6-curve irreducible Venn diagram with the DE property, derived from Grünbaum's ellipses, the original ellipses shown in Figure 4.15.

Lemma 4.2.2. *A Venn diagram that has a DE property is extendible by the heuristic described in Figure 4.1.*

Proof. We use the same labelling used to describe the conditions for the DE property. In the heuristic of Figure 4.1, we choose C_k for step 1 and v to be any vertex that is not u_k or v_k, ensuring that they are both highlighted in step 2. Then R_i, R_e, S_i, S_e are isolated faces by definition. The vertices u_m and v_m are qualifying vertices, adjacent to u_k and v_k respectively. Once both of these vertices are processed in step 11, all the faces are in V are highlighted and the additional curve produces an extended diagram. □

Theorem 4.2.3. *For every $n \geq 5$, there exist simple irreducible Venn diagrams that are extendible by Algorithm 4.1.*

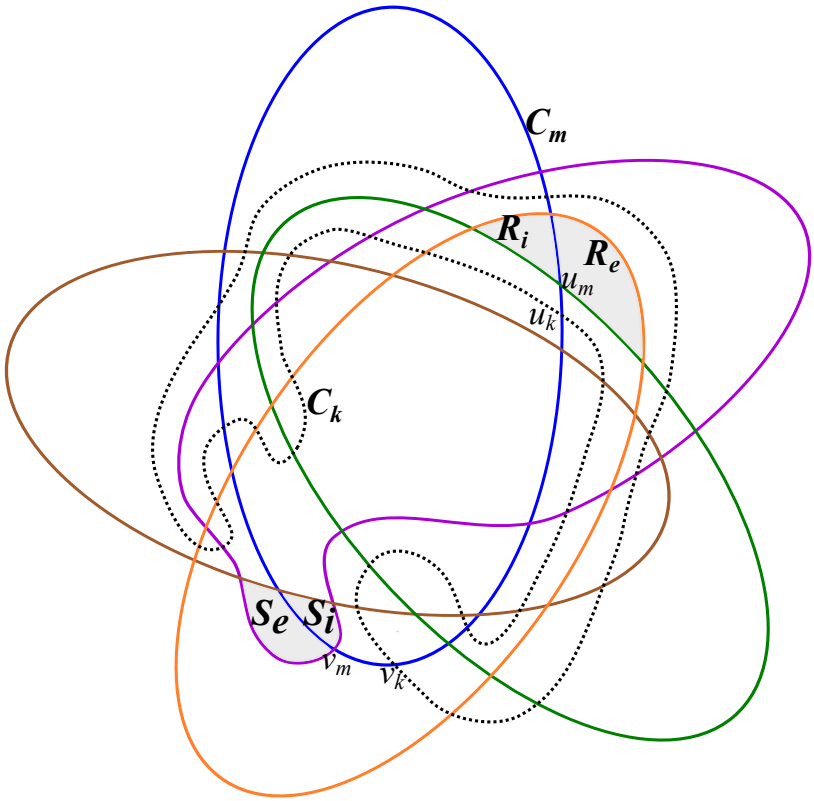

Figure 4.5: An irreducible 6-curve Venn diagram with the DE property. Curve C_k is the dashed black curve. Curve C_m is the blue curve. The faces R_i, R_e, S_i, S_e and vertices u_m, u_k, v_m, v_k are shown as described in the definition of the double extendibility property.

Proof. We use Grünbaum's inductive construction in [15] and show that each resulting Venn diagram has the DE property.

Using the same labels as in Lemma 4.2.2, we note the following facts about the regions of an n-curve Venn diagram with the DE property:

$$R_e = R_i \setminus \{m\}, S_e = S_i \setminus \{m\}, S_i = R_i \setminus \{k\}, S_e = R_e \setminus \{k\}. \tag{4.2.1}$$

Grünbaum's construction produces an irreducible Venn diagram with $n+1$ curves in the following two steps.

1. The first step modifies curve C_k to C_k^* to intersect R_i and R_e. Essentially, it takes the edge $u_k u_m$ and flips it over, preserving all the original edges incident to u_k and u_m. Figure 4.6 shows the curves before and after the edge is flipped. The original Venn diagram V is no longer a Venn diagram but an independent

set V^* by the addition of two new faces, f_i and f_e. Let f_i be interior to C_m and f_e exterior. Because they are separated by C_k^*, the region $f_i = R_i \setminus \{k\}$ and the region $f_e = R_e \setminus \{k\}$. By the identities of (4.2.1), we know specifically that f_i is equivalent to S_i and F_e is equivalent to S_e.

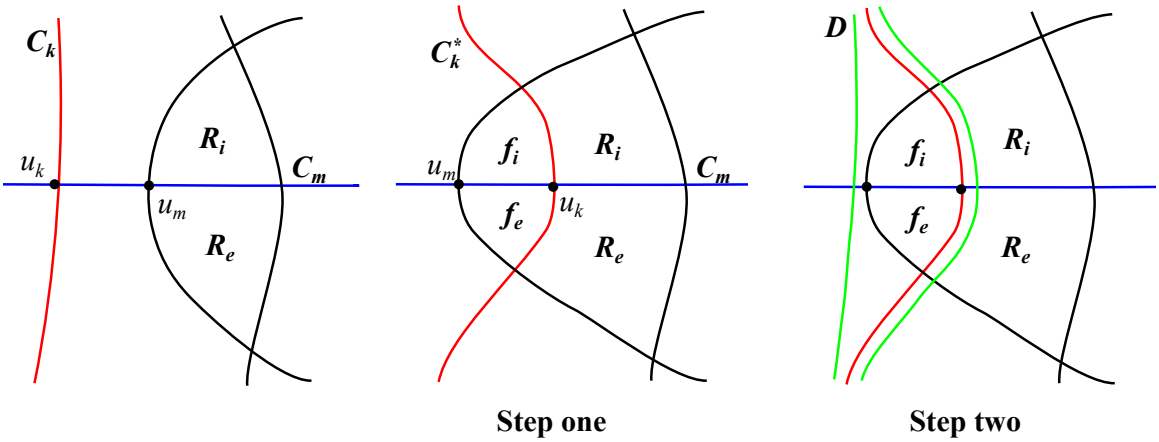

Figure 4.6: Steps one and two of the proof: Step one creates the independent set by flipping the edge $u_k u_m$ and creating two new faces. Step two adds the new curve that creates the next irreducible Venn diagram.

2. The second step adds a new curve B to V^* to form a Venn diagram with $n+1$ curves. The curve closely follows the interior and exterior of C_k^* where it is the same as C_k and crosses C_k^* on either side of a vertex w that is not v_k. Where C_k^* differs from the original C_k, B follows along the previous path of C_k on the exterior and along C_k^* on the interior. Figure 4.6 shows a segment of the new curve B around the area where C_k is altered.

The detailed proof that $V^* \cup B$ is irreducible is shown in [15]; we provide the argument that it is a Venn diagram here. The curve B, by avoiding the four faces $\{f_i, f_e, S_i, S_e\}$, separates each pair of identical faces (f_e, S_e) and (f_i, S_i), by making the f set interior and the S set exterior to B. Clearly these four faces are all unique regions in $V \cup B$.

The other $2^n - 2$ faces in V, all of which are regions, are split into unique regions interior and exterior to B in the new diagram. Splitting these regions creates $2^{n+1} - 4$ regions; Adding the four regions that do not have B on their boundary, totals 2^{n+1} unique faces, or regions, and thus $V \cup B$ is a Venn diagram.

Substituting B with C_k and f_i with R_i and f_e with R_e, all the requirements for the DE property hold for $V^* \cup B$. Therefore, by Lemma 4.2.2 this new irreducible Venn diagram of $V \cup B$ can be extended following the heuristic of Figure 4.1. By the

principle of induction, there exist simple irreducible Venn diagrams for all $n \geq 5$ for which the heuristic will produce the extending curve. □

4.3 The Venn diagrams on 5 curves

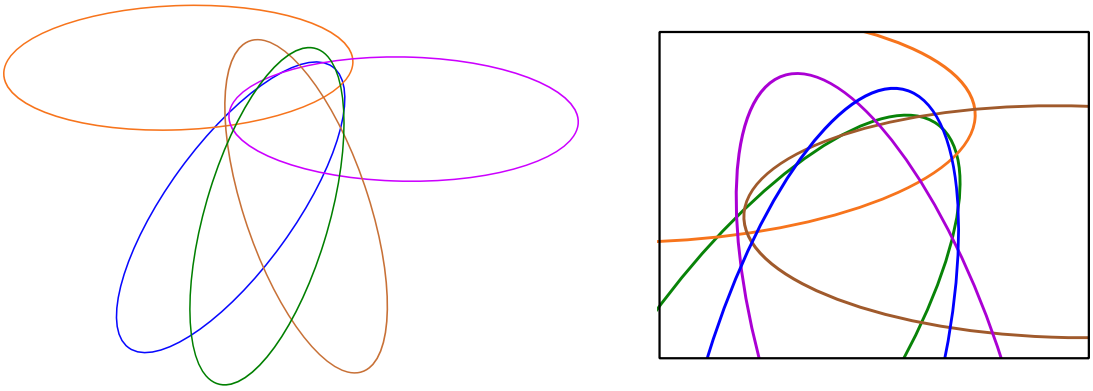

Figure 4.7: The 18th monotone embedding of the 5-curve Venn diagrams, with inset to show detail.

Table 4.1: The dimensions of the ellipses, illustrated in Figure 4.7. The units are in pixels, where the lower left point (x, y) of the diagram is $(0, 0)$.

curve	x	y	width	height
orange	0.006	38.809	100.801	280.8
blue	38.558	77.411	274.763	117.966
green	51.473	106.546	229.717	191.495
brown	44.010	46.302	250.504	162.287
purple	55.902	0.007	217.741	202.776

All 20 classes of Venn diagrams on five curves were identified and catalogued by Hamburger et al. in [18] and [10]. We verify the existence of the unique diagrams, mapping each to its original figure in the two papers. In [10], 17 non-isomorphic convex embeddings on the plane are derived from 11 classes of convex diagrams. We discovered one more, correcting the previous total to 18. Specifically, the right lower drawing of Figure 2 in [10] has another monotone, and therefore convex, embedding that results when the curve with 8 intersections has its interior flipped to its exterior. This diagram can also be drawn with congruent ellipses, as shown in Figure 4.7. The dimension of each of the bounding rectangles for the ellipses is given in Table 4.1.

We categorize all the Venn diagrams (by class) in the following table, with the following information:

- The figure number of the Venn diagram.
- Whether the diagram is reducible.
- The number of unique monotone embeddings for the diagram.
- The sequence of the number of intersection points of each curve, from largest number to smallest.
- The listing of the face distribution as a tuple k_3, k_4, k_5, where k_i is the number of faces bounded by i curve segments.

Table 4.2: All Venn diagrams on five curves

figure	reducible?	monotone embeddings	crossing distribution	face distribution
4.8	no	0	$14, 14, 12, 12, 8$	$14, 12, 6$
4.9	no	0	$14, 14, 12, 12, 8$	$12, 16, 4$
4.10	no	0	$14, 14, 12, 12, 8$	$14, 12, 6$
4.11	no	0	$14, 12, 12, 12, 10$	$14, 12, 6$
4.12	no	1	$14, 14, 12, 10, 10$	$12, 16, 4$
4.13	no	1	$14, 14, 12, 10, 10$	$15, 11, 6$
4.14	no	1	$14, 14, 12, 12, 8$	$16, 8, 8$
4.15	no	1	$12, 12, 12, 12, 12$	$10, 20, 2$
4.16	no	2	$14, 14, 12, 12, 8$	$14, 12, 6$
4.17	yes	0	$16, 16, 10, 10, 8$	$10, 20, 2$
4.18	yes	0	$16, 14, 12, 10, 8$	$12, 16, 4$
4.19	yes	0	$16, 12, 12, 12, 8$	$12, 16, 4$
4.20	yes	0	$16, 14, 10, 10, 10$	$14, 12, 6$
4.21	yes	0	$16, 14, 10, 10, 10$	$10, 20, 2$
4.22	yes	1	$16, 16, 12, 8, 8$	$16, 8, 8$
4.23	yes	1	$16, 12, 12, 12, 8$	$12, 16, 4$
4.24	yes	2	$16, 14, 14, 8, 8$	$16, 8, 8$
4.25	yes	2	$16, 14, 12, 10, 8$	$14, 12, 6$
4.26	yes	2	$16, 12, 12, 10, 10$	$14, 12, 6$
4.27	yes	4	$16, 14, 12, 10, 8$	$14, 12, 6$

Each of the monotone Venn diagrams is drawn using a monotone embedding. The coloured stars in the faces indicate that a monotone embedding is also possible with this identified region as the outer region. Same-coloured stars indicate the embeddings are isomorphic.

Although each of the monotone Venn diagrams can be drawn with convex curves, this is not necessarily shown when the convexity interferes with the relative sizes of the regions. It is the author's preference that a smooth curve and reasonably proportioned regions are more pleasing to the eye.

Each non-monotone Venn diagram is drawn exposed. Each Venn diagram has an easy extension to a 6-curve Venn diagram, indicated by the dashed curve following the green Venn curve as the base curve.

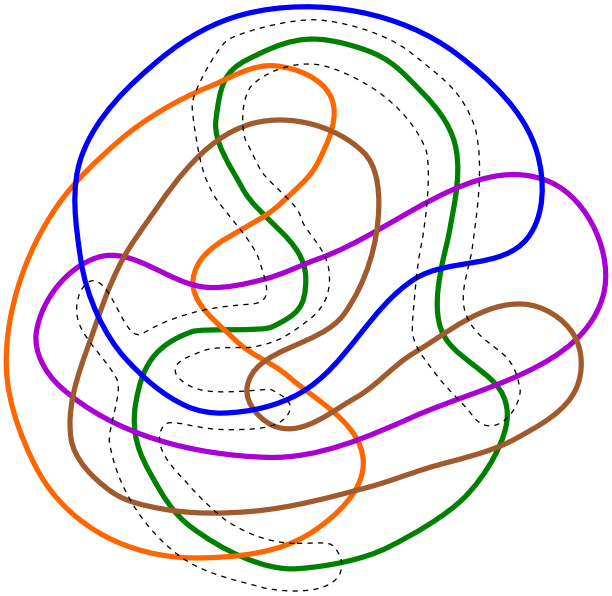

Figure 4.8: I_1: An irreducible, non-monotone Venn diagram on 5 curves, Figure 1(i) from [10].

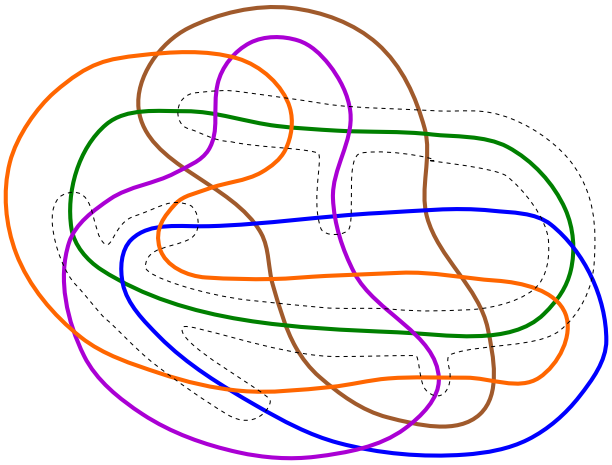

Figure 4.9: I_2: An irreducible, non-monotone Venn diagram on 5 curves, Figure 1(iv) from [10].

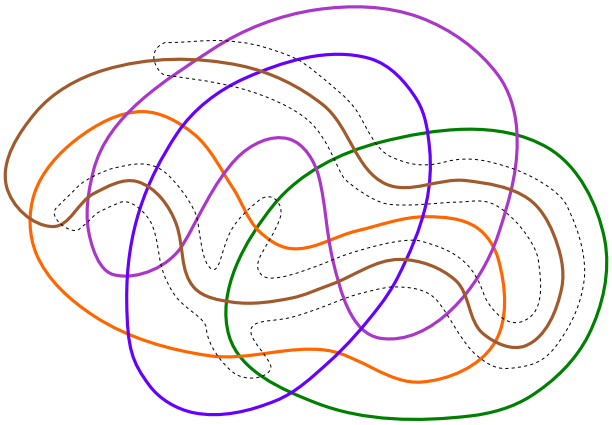

Figure 4.10: I_3: An irreducible, non-monotone Venn diagram on 5 curves, Figure 1(ii) from [10].

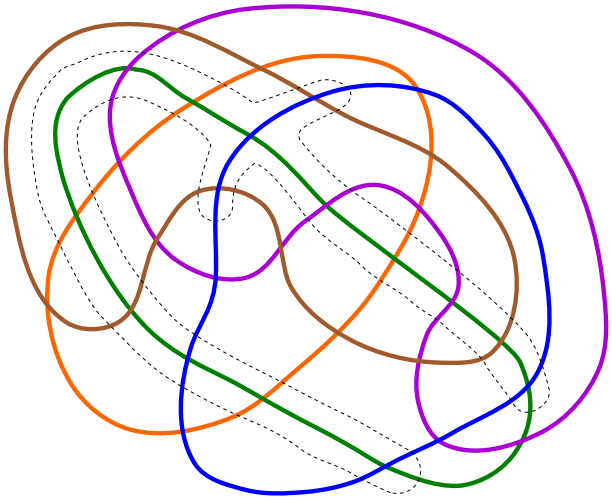

Figure 4.11: I_4: An irreducible, non-monotone Venn diagram on 5 curves, Figure 1(iii) from [10].

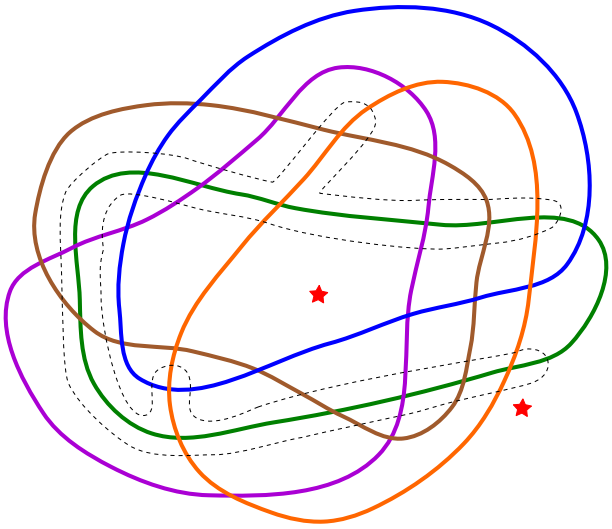

Figure 4.12: I_5: An irreducible, monotone Venn diagram on 5 curves, one unique monotone embedding, Figure 1 from [10].

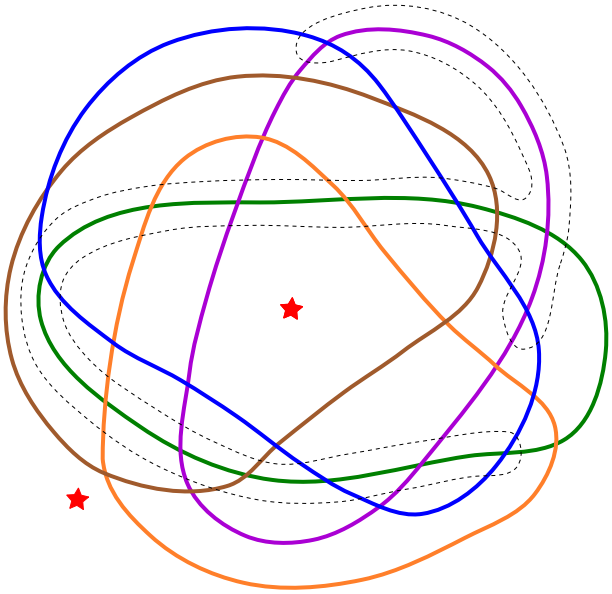

Figure 4.13: I_6: An irreducible, monotone Venn diagram on 5 curves, one unique monotone embedding, Figure 2(ii) from [10].

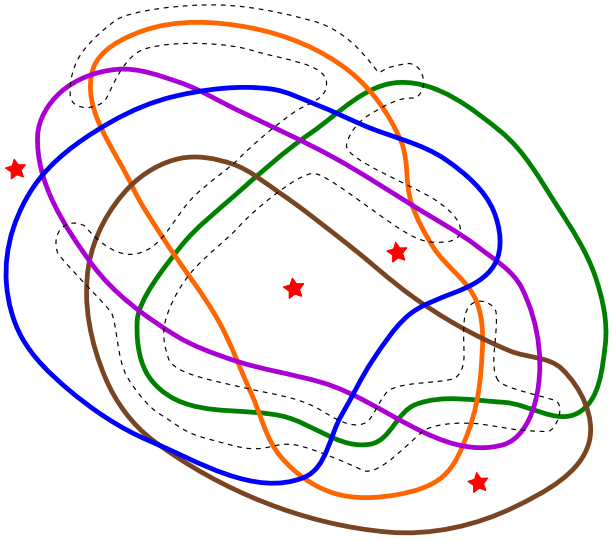

Figure 4.14: I_7: An irreducible, monotone Venn diagram on 5 curves, one unique monotone embedding, Figure 2(i) from [10].

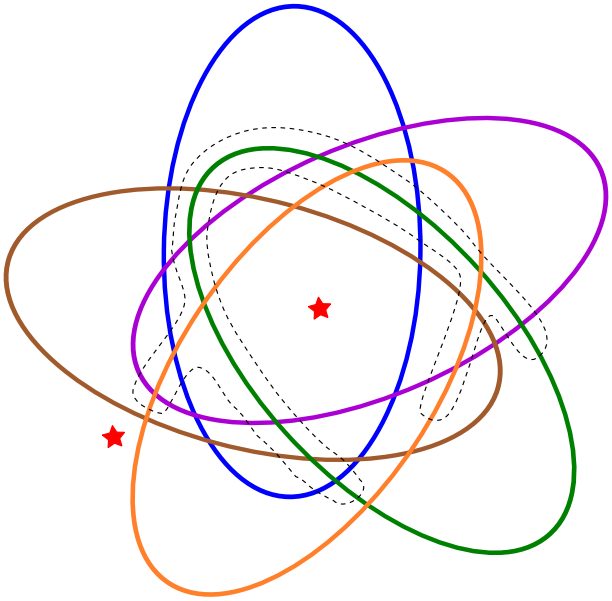

Figure 4.15: I_8: Grünbaum's ellipses, one unique monotone embedding.

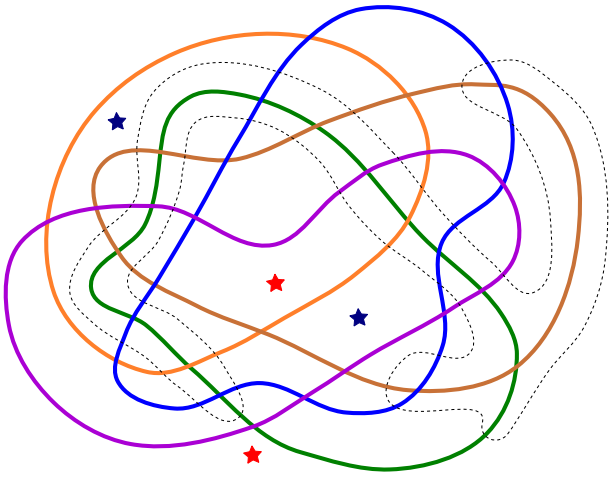

Figure 4.16: I_9: An irreducible, monotone Venn diagram on 5 curves, Figure 2 (unlabelled) from [10]. This is the diagram that yields the additional monotone embedding that can be drawn using congruent ellipses.

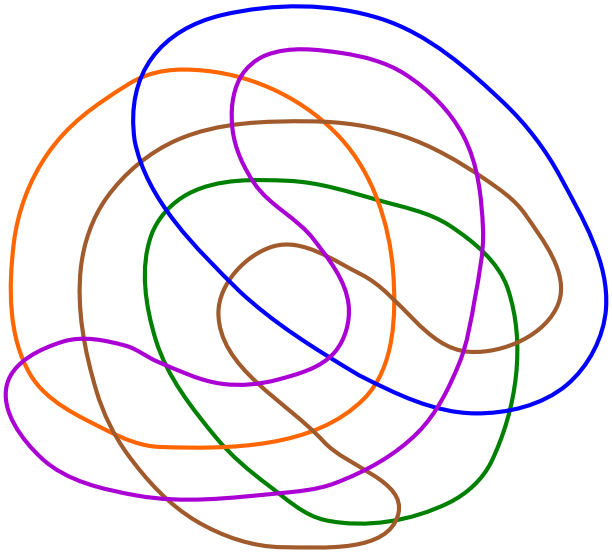

Figure 4.17: R_1: A reducible, non-monotone Venn diagram on 5 curves, Figure 4 (top right) from [18].

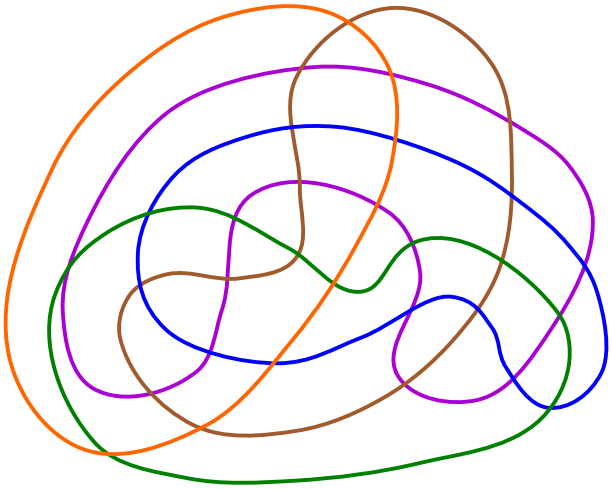

Figure 4.18: R_2: A reducible, non-monotone Venn diagram on 5 curves, Figure 4(second row right) from [18].

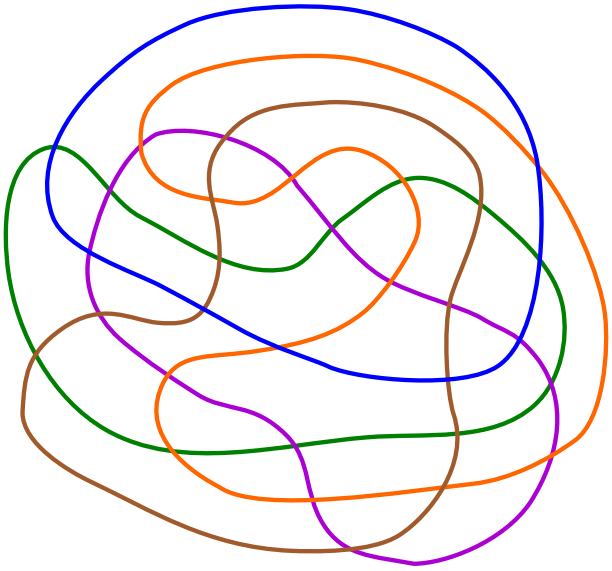

Figure 4.19: R_3: A reducible, non-monotone Venn diagram on 5 curves, Figure 4(second row left) from [18].

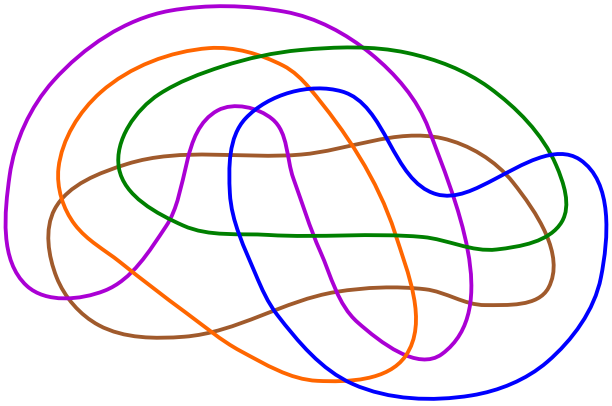

Figure 4.20: R_4: A reducible, non-monotone Venn diagram on 5 curves, Figure 4(bottom) from [18].

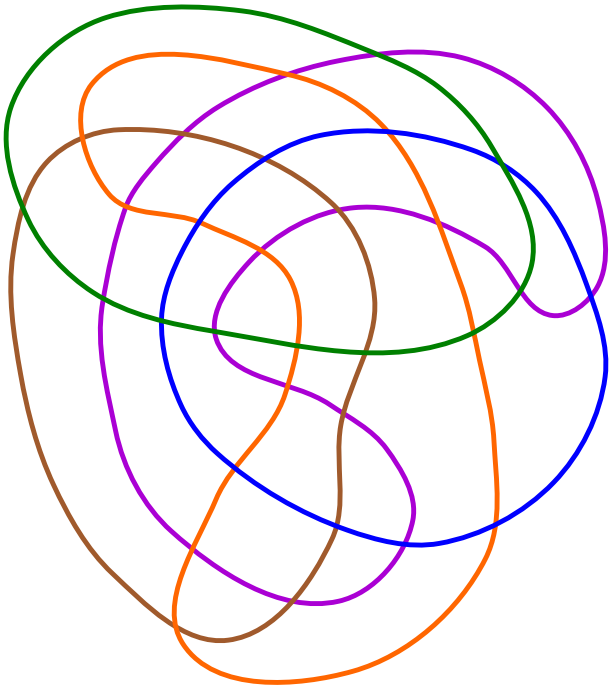

Figure 4.21: R_5: A reducible, non-monotone Venn diagram on 5 curves, Figure 4(top left) from [18].

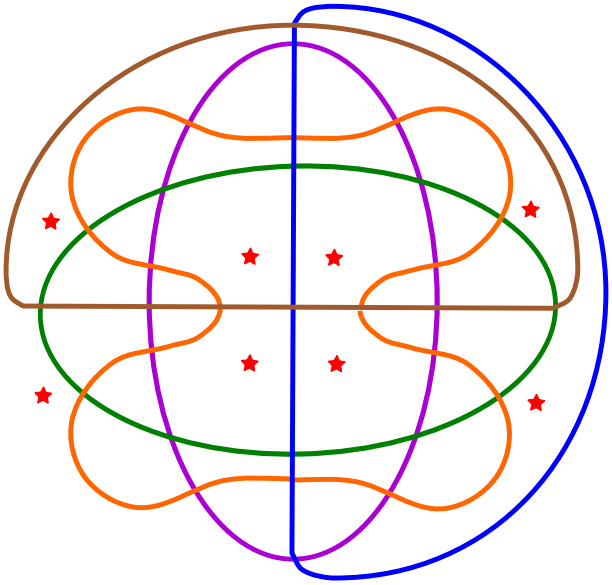

Figure 4.22: R_6: Edward's construction of a reducible, monotone Venn diagram on 5 curves, one unique monotone embedding, Figure 3(Diamond) from [18].

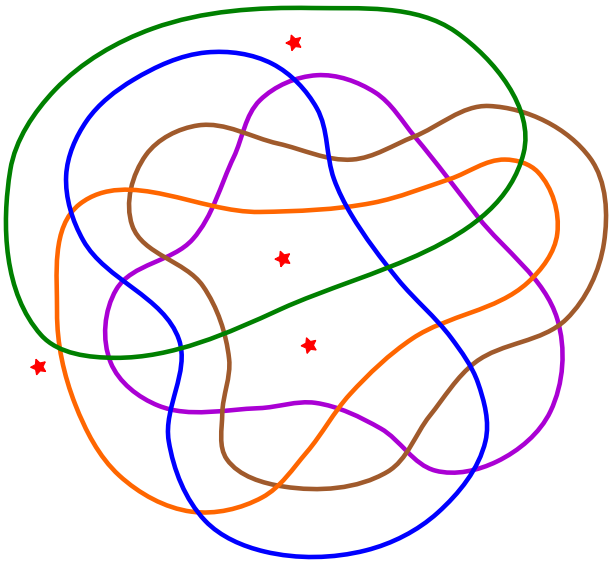

Figure 4.23: R_7: A reducible, monotone Venn diagram on 5 curves, one unique monotone embedding, Figure 3(E) from [18].

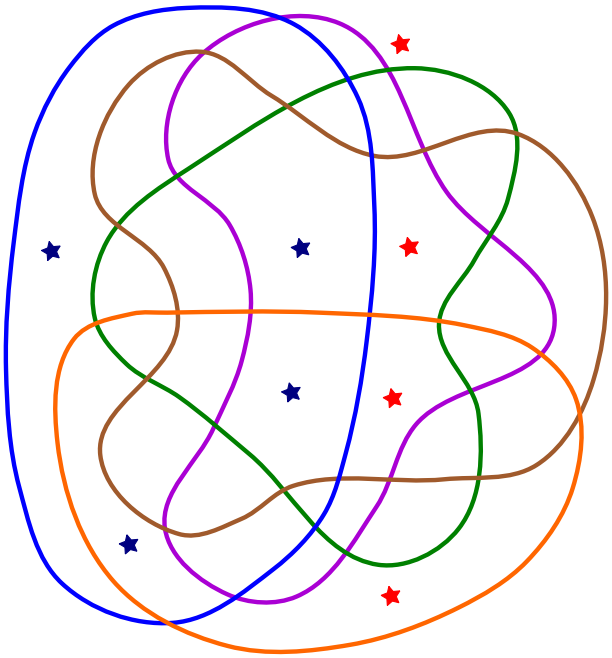

Figure 4.24: R_8: A reducible, monotone Venn diagram on 5 curves, two unique monotone embeddings, Figure 3(Goal post) from [18].

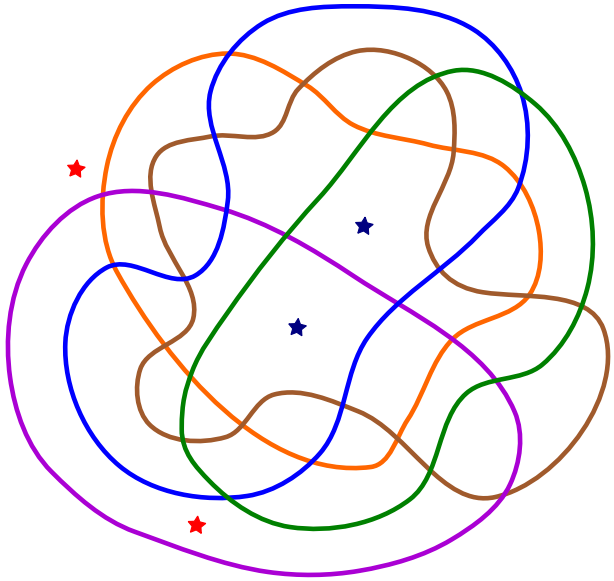

Figure 4.25: R_9: A reducible, monotone Venn diagram on 5 curves, two unique monotone embeddings, Figure 3(C) from [18].

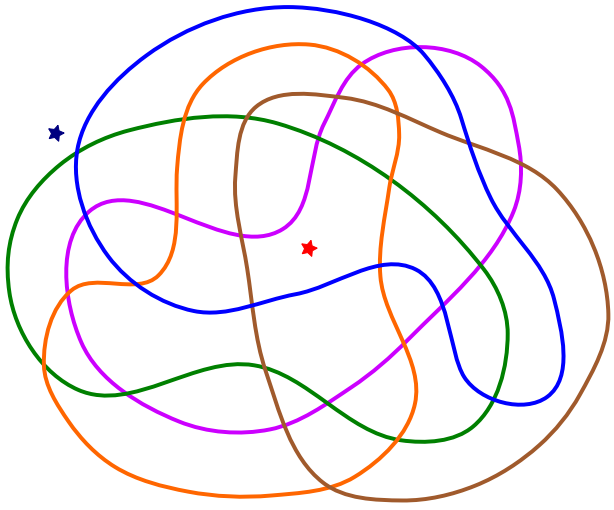

Figure 4.26: R_{10}: A reducible, monotone Venn diagram on 5 curves, two unique monotone embeddings, Figure 3(The clown) from [18].

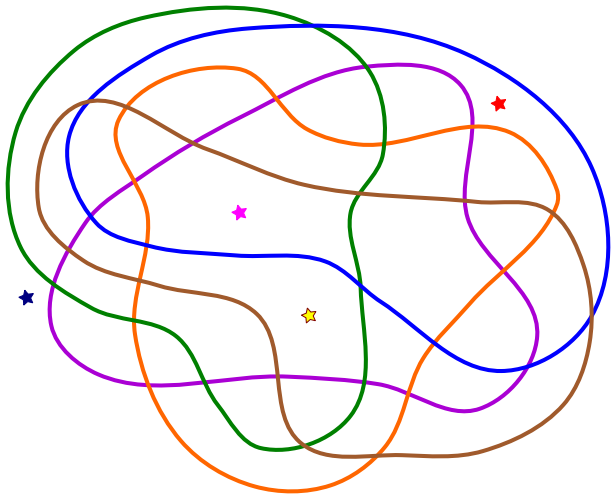

Figure 4.27: R_{11}: A reducible, monotone Venn diagram on 5 curves, four unique monotone embeddings, Figure 3(Chair) from [18].

4.3.1 Venn diagrams, face-balanced curves and set theory

Not only does each Jordan curve in a Venn diagram divide the plane into two sets of 2^{n-1} faces, but each set of two curves divides the plane into four sets of 2^{n-2} faces. This continues so that each set of $k \leq n$ curves divides the plane into 2^k sets of 2^{n-k} faces. Venn diagrams are often discussed in terms of set theory so it seems reasonable to look at the "half" sets themselves. We look at this in the next chapter.

Chapter 5

The Half-Set System

In this chapter, we examine the properties of a Venn diagram relating to the division of the number of faces between the curves. We define this in terms of a system of subsets of integers $\{1, 2, \ldots, n\}$. In Chapter 6, we use this set system to extend a Venn diagram where every connected region is represented as a single unit square.

Definition 5.0.1. A *half-set system* on 2^n elements, or an n-HSS, is a collection $\mathbf{S} = \{S_1, S_2, \ldots, S_n\}$ of subsets of $\{1, 2, \ldots, 2^n\}$ with the property that for any non-empty subset $A \subseteq \{1, 2, \ldots, n\}$

$$\left| \bigcap_{i \in A} S_i \right| = 2^{n-|A|}.$$

Note that each subset S_i has cardinality 2^{n-1}, namely it is a "half-set".

The following theorem provides an alternate way of checking for the HSS property.

Theorem 5.0.2. *Let* $\mathbf{S} = \{S_1, S_2, \ldots, S_n\}$ *be a collection of subsets of* $\{1, 2, \ldots, 2^n\}$. *Then* $\mathbf{S}$ *is an HSS if and only if there is a bijection:* $f : \mathcal{P}(\{1, 2, \ldots, n\}) \longrightarrow \{1, 2, \ldots, 2^n\}$ *such that for all* $B \subseteq \{1, 2, \ldots, n\}$,

$$\left(\bigcap_{i \in B} S_i \right) \cap \left(\bigcap_{i \notin B} \overline{S}_i \right) = \{f(B)\}. \tag{5.0.1}$$

The proof of this Theorem is somewhat long and technical, so it can be found in Section 5.1.1, at the end of this chapter. Equation (5.0.1) is analogous to Grünbaum's definition, stated in Section 1.2; each $X_1 \cap X_2 \cap \cdots \cap X_n$, where X_i is the bounded

interior of curve C_i if $i \in B$ and the unbounded exterior of curve C_i if $i \notin B$ is a single region.

The following lemma demonstrates that under certain conditions, an HSS is *expandable*, by the addition of more elements.

Lemma 5.0.3. *Let S be a set of 2^n elements, for some positive integer n. Let $X(S)$ be an expanded set of 2^{n+c} elements such that for every $e_i \in S$, there exists a corresponding set $X(e_i)$ of 2^c elements, where c is a positive integer and for every $i \neq j$ and $1 \leq i, j \leq 2^n$, $X(e_i) \cap X(e_j) = \emptyset$. Let $\mathbf{H} = \{H_1, H_2, \ldots, H_n\}$ be an HSS on S. Let $X(\mathbf{H})$ be a set of subsets $\{X(H_1), X(H_2), \ldots, X(H_n)\}$ on $X(S)$ where $X(H_j)$ has the following property:*

$$\text{If } e_i \in S_j, \text{ then } X(e_i) \subseteq X(H_j). \tag{5.0.2}$$

For every $i \in \{1, 2, \ldots, 2^c\}$, let $\mathbf{M}_i = \{M_{i1}, M_{i2}, \ldots, M_{ic}\}$ be an HSS on $X(e_i)$. Let $\mathbf{E}$ be a set of subsets on $X(S)$, $\mathbf{E} = \{E_1, E_2, \ldots, E_c\}$ with the conditions that for all $k \in \{1, 2, \ldots, c\}$,

$$E_k = \bigcup_{1 \leq i \leq c} M_{ik}.$$

Then $X(\mathbf{H}) \cup \mathbf{E}$ is an HSS on $X(S)$.

See Section 5.1.2 for the proof.

Consider the regions of a Venn diagram V as the elements of an HSS. Lemma 5.0.3 tells us that if we add curves or curve segments to each region that expands it into 2^c new regions and if these new regions are the elements of an HSS, then the resulting diagram is an HSS. This is obviously not a necessary condition for extending a Venn diagram; imagine if we added the same little 3-curve Venn diagram in each region of V. If the three curves in each region are labelled C_{n+1}, C_{n+2} and C_{n+3}, the result is an HSS, but most definitely not a Venn diagram. In the following chapter, we show that the c curves must be connected. We also show that Lemma 5.0.3 provides one of the necessary conditions to expand Venn diagrams where every region is a single unit square whose points lie on the integer lattice.

5.1 Proofs

5.1.1 Proof of Theorem 5.0.2

We begin by demonstrating another formula associated with the HHS property that is presented as part of the proof of this theorem.

Lemma 5.1.1. *Given an HSS,* $\mathbf{S} = \{S_1, S_2, \ldots, S_n\}$ *and any non-empty subset* $A \subseteq \{1, 2, \ldots, n\}$*, then for any* $X \subseteq A$ *and* $Y = A \setminus X$*,*

$$\left| \left(\bigcap_{i \in X} S_i \right) \cap \left(\bigcap_{i \in Y} \overline{S_i} \right) \right| = 2^{n-|A|}. \tag{5.1.1}$$

Proof. Our proof is by induction on increasing values of $|Y|$ for any fixed A where $|A| \geq 1$. We assume that the intersection of a family of sets indexed over the empty set is the universe. Thus, when $Y = \emptyset$, $X = A$ and

$$\left| \left(\bigcap_{i \in X} S_i \right) \right| = 2^{n-|X|} = 2^{n-|A|},$$

which is true by Definition 5.0.1. If A contains only one element e, and $X = \emptyset$, we note that S_e and $\overline{S_e}$ are disjoint sets, so $|\overline{S_e}| = 2^n - 2^{n-1} = 2^{n-1}$. We have proved the base case for $|Y| = 0$ and the case when $|A| = 1$ and $X = \emptyset$.

Assume that (5.1.1) is true for all $0 \leq |Y| < |A| = n$, for some fixed value $n > 1$. Let $e \in X$, let $B = A \setminus \{e\}$ and let

$$Z = \left(\bigcap_{i \in X \setminus \{e\}} S_i \right) \cap \left(\bigcap_{i \in Y} \overline{S_i} \right).$$

By the induction hypothesis, $|Z \cap S_e| = 2^{n-|A|}$ and because B is non-empty, $|Z| = 2^{n-|B|} = 2^{n-(|A|-1)}$.

Moving e from X into Y, we have

$$\left(\bigcap_{i \in X \setminus \{e\}} S_i \right) \bigcap \left(\bigcap_{i \in Y \cup \{e\}} \overline{S_i} \right) = Z \bigcap \overline{S_e}.$$

By basic set principles, $Z = (Z \cap S_e) \cup (Z \cap \overline{S_e})$, and $|Z| = |Z \cap S_e| + |Z \cap \overline{S_e}|$, so $|Z \cap \overline{S_e}| = 2^{n-(|A|-1)} - 2^{n-|A|} = 2^{n-|A|}$. □

The following lemma demonstrates that Equation (5.1.1) above is also sufficient to define the HSS property.

Lemma 5.1.2. *Let* $\mathbf{S} = \{S_1, S_2, S_3, \dots S_n\}$ *be a collection of subsets on* $\{1, 2, \dots, 2^n\}$. *If* $\mathbf{S}$ *has the property that for any non-empty set* A *and any* $X \subseteq A$ *and* $Y = A \setminus B$ *where the following is true:*

$$\left| \left(\bigcap_{i \in X} S_i \right) \cap \left(\bigcap_{i \in Y} \overline{S_i} \right) \right| = 2^{n-|A|)},$$

then $\mathbf{S}$ *is an HSS.*

Proof. Let $\mathbf{T}$ be a set of subsets $\{T_1, T_2, \dots, T_n\}$ where

$$T_i = \begin{cases} S_i & \text{if } i \in X \\ \overline{S_i} & \text{if } i \in Y. \end{cases}$$

Then

$$\begin{aligned} \left| \bigcap_{i \in A} T_i \right| &= \left| \left(\bigcap_{i \in X} S_i \right) \cap \left(\bigcap_{i \in Y} \overline{S_i} \right) \right| \\ &= 2^{n-|A|} \end{aligned}$$

and $\mathbf{T}$ is an HSS by Definition 5.0.1. In that case,

$$\left| \left(\bigcap_{i \in X} T_i \right) \cap \left(\bigcap_{i \in Y} \overline{T_i} \right) \right| = 2^{n-|A|} \text{ by Lemma 5.1.1.}$$

However,

$$\left(\bigcap_{i \in X} T_i \right) \cap \left(\bigcap_{i \in Y} \overline{T_i} \right) = \left(\bigcap_{i \in X} S_i \right) \cap \left(\bigcap_{i \in Y} S_i \right) = \bigcap_{i \in A} S_i.$$

By Definition 5.0.1, $\mathbf{S}$ is an HSS. □

Copy of Theorem. *Let* $\mathbf{S} = \{S_1, S_2, \dots, S_n\}$ *be a collection of subsets of* $\{1, 2, \dots, 2^n\}$. *Then* $\mathbf{S}$ *is an HSS if and only if there is a bijection:* $f : \mathcal{P}(\{1, 2, \dots, n\}) \longrightarrow \{1, 2, \dots, 2^n\}$ *such that for all* $B \subseteq \{1, 2, \dots, n\}$,

$$\left(\bigcap_{i \in B} S_i \right) \cap \left(\bigcap_{i \notin B} \overline{S}_i \right) = \{f(B)\}.$$

Proof. ($\Rightarrow$) Let $B, C \subseteq \{1, 2, \ldots, n\}$ and $e \in B \setminus C$. By Lemma 5.1.1,

$$|f(B)| = \left| \left(\bigcap_{i \in B} S_i \right) \bigcap \left(\bigcap_{i \notin B} \overline{S_i} \right) \right| = 1,$$

and $f(C)$ is defined similarly. However, $F(B) \in S_e$ and $f(C) \in \overline{S_e}$. Thus $f(B) \neq F(C)$ and f is one-to-one. Because $|\mathcal{P}\{1, 2, \ldots, n\}| = |\{1, 2, \ldots, 2^n\}| = 2^n$, f is onto.

($\Leftarrow$) Given $\mathbf{S} = \{S_1, S_2, \ldots, S_n\}$ is a set of subsets on $\{1, 2, \ldots, n\}$ and for all $B \subseteq \{1, 2, \ldots n\}$, there exists a bijection from B to $f(B) \in \{1, 2, \ldots, 2^n\}$ such that

$$\left(\bigcap_{i \in B} S_i \right) \cap \left(\bigcap_{i \notin B} \overline{S_i} \right) = \{f(B)\}.$$

Let $A \subseteq \{1, 2, \ldots, n\}$ and $B \subseteq A$. Let

$$Z = \left(\bigcap_{i \in B} S_i \right) \cap \left(\bigcap_{i \in A \setminus B} \overline{S_i} \right). \tag{5.1.2}$$

We use induction on decreasing values of $|A|$ to show that:

$$|Z| = 2^{n-|A|} \text{ for all } 0 < |A| \leq n. \tag{5.1.3}$$

For the base case, $|A| = n$. Then for all $B \subseteq A$, $Z = \{m\}$ and $|Z| = 1 = 2^0 = 2^{n-|A|}$. Suppose (5.1.3) is true for some $1 < |A| \leq n$. Choose a non-empty subset $C \subseteq A$. Let $e \in C$ and $B = C \setminus \{e\}$. Let

$$X = \left(\bigcap_{i \in C} S_i \right) \cap \left(\bigcap_{i \in A \setminus C} \overline{S_i} \right)$$

and

$$Y = \left(\bigcap_{i \in B} S_i \right) \cap \left(\bigcap_{i \in A \setminus B} \overline{S_i} \right).$$

By the induction hypothesis, $|X| = |Y| = 2^{n-|A|}$. Let $Z' = X \cup Y$. Because X and Y

are disjoint sets, $|Z'| = |X \cup Y| = |X| + |Y| = 2 \times 2^{n-|A|} = 2^{n-(|A|-1)}$. Furthermore:

$$\begin{aligned} Z' = &\left[\left(\bigcap_{i\in B} S_i\right) \cap \left(\bigcap_{i\in A\setminus C} \overline{S_i}\right) \cap S_e\right] \cup \\ &\left[\left(\bigcap_{i\in B} S_i\right) \cap \left(\bigcap_{i\in A\setminus C} \overline{S_i}\right) \cap \overline{S_e}\right] \\ = &\left(\bigcap_{i\in B} S_i\right) \cap \left(\bigcap_{i\in A\setminus C} \overline{S_i}\right). \end{aligned}$$

Letting $A' = A \setminus \{e\}$ and noting that $B \subseteq A'$ was arbitrarily chosen,

$$|Z'| = \left|\left(\bigcap_{i\in B} S_i\right) \cap \left(\bigcap_{i\in A'\setminus B} \overline{S_i}\right)\right| = 2^{n-|A\setminus\{e\}|} = 2^{n-|A'|}$$

and (5.1.3) is shown for all $|A|$. By Lemma 5.1.2, $\mathbf{S}$ is an HSS. □

5.1.2 Proof of Lemma 5.0.3

We begin with some important facts that are used in the proof of this lemma. Both facts use the following assumptions: Let $\{X_1, X_2, \ldots X_n\}$ be set of pairwise disjoint sets. For any $m > 0$, let $\mathbf{X} = \{X_{ij} \mid 1 \leq i \leq n, 1 \leq j \leq m, X_{ij} \subseteq X_i\}$. So $X_{ij} \cap X_{uv} = \emptyset$ when $i \neq u$.

Fact 5.1.3. *For any non-empty sets $A \subseteq \{1, 2, \ldots, n\}$ and $B \subseteq \{1, 2, \ldots, m\}$, the following statement is true:*

$$\bigcap_{j\in B}\left(\bigcup_{i\in A} X_{ij}\right) = \bigcup_{i\in A}\left(\bigcap_{j\in B} X_{ij}\right). \tag{5.1.4}$$

Proof. Without loss of generality, we can permute the X_{ij} sets to match the indices in A and B, so $A = \{1, 2, \ldots k\}$ and $B = \{1, 2, \ldots, q\}$. Consider the lefthand side of (5.1.4) which, when expanded into the union of intersections, is similar to the Cartesian products in that each intersection chooses one element from each set of

unions. Letting the intersection operators be implied, we get

$$\bigcap_{j=1}^{q}\left(\bigcup_{i=1}^{k} X_{ij}\right) = \bigcup_{\substack{j=1,\ldots,q \\ i_j=1,\ldots,k}} X_{i_1 1} X_{i_2 2} X_{i_3 3} \cdots X_{i_q q}$$

for a total of q^k unions. Because $X_{ij} \cap X_{uv} = \emptyset$ when $i \neq u$, we can delete all the union terms where the i indices are not equal, leaving

$$(X_{11}X_{12}\ldots X_{1q}) \cup (X_{21}X_{22}\ldots X_{2q}) \cup \ldots \cup (X_{k1}X_{k2}\ldots X_{kq}) = \bigcup_{i=1}^{k}\left(\bigcap_{j=1}^{q} X_{ij}\right).$$

□

Fact 5.1.4. *Let $Y_i \subseteq X_i$. Let $Z = \cup_{i=1}^{n} Y_i$. Then the complement of Z is*

$$\overline{Z} = \bigcup_{i=1}^{n} \overline{Y_i},$$

where $\overline{Y_i} = X_j \setminus Y_i$.

Proof. Let $\overline{Z} = \mathbb{U} \setminus Z$, where $\mathbb{U} = \bigcup_{i=1}^{n} X_i$. Let $x \in \overline{Z}$. Then there is some i such that $x \in X_i$ and $x \notin Y_i$. Thus $x \in \cup_i (X_i \setminus Y_i)$ and $\overline{Z} \subseteq \cup_i (X_i \setminus Y_i)$. On the other hand, suppose $x \in \cup_i (X_i \setminus Y_i)$. Then there is some i such that $x \in X_i$ but $x \notin Y_i$. Because i is unique and the X_is are disjoint, $x \in \mathbb{U}$ and $x \notin Z$. Thus $x \in \overline{Z}$ and $\cup_i (X_i \setminus Y_i) \subseteq \overline{Z}$. □

Copy of Lemma. *Let S be a set of 2^n elements, for some positive integer n. Let $X(S)$ be an expanded set of 2^{n+c} elements such that for every $e_i \in S$, there exists a corresponding set $X(e_i)$ of 2^c elements, where c is a positive integer and for every $i \neq j$ and $1 \leq i, j \leq 2^n$, $X(e_i) \cap X(e_j) = \emptyset$. Let $\mathbf{H} = \{H_1, H_2, \ldots, H_n\}$ be an HSS on S. Let $X(\mathbf{H})$ be a set of subsets, $\{X(H_1), X(H_2), \ldots, X(H_n)\}$ on $X(S)$ where $X(H_j)$ has the following property:*

$$\textit{If } e_i \in S_j, \textit{ then } X(e_i) \subseteq X(H_j).$$

Let $\mathbf{M}_i = \{M_{i1}, M_{i2}, \ldots, M_{ic}\}$ be an HSS on $X(e_i)$. Let $\mathbf{E}$ be a set of subsets on

$X(S)$, $\mathbf{E} = \{E_1, E_2, \ldots, E_c\}$ *with the conditions that for all* $k \in \{1, 2, \ldots, c\}$,

$$E_k = \bigcup_{1 \leq i \leq c} M_{ik}. \tag{5.1.5}$$

Then $X(\mathbf{H}) \cup \mathbf{E}$ *is an HSS on* $X(S)$.

Proof. We begin by identifying the following universal sets:

$$\begin{aligned} \mathbb{U}_n &= \{1, 2, \ldots, n\}, \\ \mathbb{U}_c &= \{1, 2, \ldots, c\}, \\ \mathbb{U} &= \{1, 2, \ldots, n + c\}. \end{aligned}$$

Choose an element $e_u \in S$, where $1 \leq u \leq 2^n$. Since $\mathbf{H}$ is an HSS over S, by Theorem 5.0.2 there is a subset $A \subseteq \mathbb{U}_n$ such that

$$\{e_u\} = \left(\bigcap_{k \in A} H_k \right) \bigcap \left(\bigcap_{k \in \mathbb{U}_n \setminus A} \overline{H_k} \right). \tag{5.1.6}$$

Then, by applying the definition of $X(\mathbf{H})$:

$$X(e_u) = \left(\bigcap_{k \in A} X(H_k) \right) \bigcap \left(\bigcap_{k \in \mathbb{U}_n \setminus A} \overline{X(H_k)} \right). \tag{5.1.7}$$

Choose an element $e \in X(e_u) \subset X(S)$. Because $\mathbf{M}_u$ is an HSS on $X(e_u)$, by Theorem 5.0.2 there is a subset $B \subseteq \mathbb{U}_c$ such that

$$\{e\} = \left(\bigcap_{k \in B} M_{uk} \right) \bigcap \left(\bigcap_{k \in \mathbb{U}_c \setminus B} \overline{M_{uk}} \right). \tag{5.1.8}$$

Let $Y \subseteq X(S)$ be defined by the following equation:

$$Y = \left(\bigcap_{k \in B} E_k \right) \bigcap \left(\bigcap_{k \in \mathbb{U}_c \setminus B} \overline{E_k} \right). \tag{5.1.9}$$

By the definition of E_k in (5.1.5) and $\overline{E_k}$ from Fact 5.1.4, we see

$$\begin{aligned} Y &= \left[\bigcap_{k\in B}\left(\bigcup_{1\le i\le 2^n} M_{ik}\right)\right]\bigcap\left[\bigcap_{k\in \mathbb{U}_c\setminus B}\left(\bigcup_{1\le i\le 2^n}\overline{M_{ik}}\right)\right] \\ &= \bigcup_{1\le i\le 2^n}\left[\left(\bigcap_{k\in B} M_{ik}\right)\bigcap\left(\bigcap_{k\in \mathbb{U}_c\setminus B}\overline{M_{ik}}\right)\right], \end{aligned}$$

by Fact 5.1.3 and the Distributive Principle. From Equation (5.1.8), we see that when $i = u$, e is the only element from $X(e_u)$ in Y. So

$$\{e\} = X(e_u)\cap Y. \tag{5.1.10}$$

Finally, let $C = A\cup\{k+n \mid k\in B\}$ and $\mathbf{T} = X(\mathbf{S})\cup\mathbf{E}$, where $T_j = X(S_j)$ and $T_{k+n} = E_k$. Then by Equations (5.1.10), (5.1.7) and (5.1.9):

$$\begin{aligned} \{e\} &= X(e_u)\cap Y \\ &= \left(\bigcap_{k\in A} X(H_k)\right)\bigcap\left(\bigcap_{k\in\mathbb{U}_n\setminus A}\overline{X(H_k)}\right)\bigcap\left(\bigcap_{k\in B} E_k\right)\bigcap\left(\bigcap_{k\in\mathbb{U}_c\setminus B}\overline{E_k}\right) \\ &= \left(\bigcap_{j\in A} T_j\right)\bigcap\left(\bigcap_{j\in\mathbb{U}_n\setminus A}\overline{S_j}\right)\bigcap\left(\bigcap_{k\in B} S_{k+n}\right)\bigcap\left(\bigcap_{k\in\mathbb{U}_c\setminus B}\overline{S_{k+n}}\right) \\ &= \left(\bigcap_{i\in C} T_i\right)\bigcap\left(\bigcap_{i\in\mathbb{U}\setminus C}\overline{T_i}\right). \end{aligned}$$

By Theorem 5.0.2, $\mathbf{T} = X(\mathbf{S})\cup E$ is an HSS. □

Chapter 6

Minimum Area Venn Diagrams

This chapter investigates Venn diagrams in which each region is a single unit square; the research was joint work with Matthew Klimesh and published in [5].

Venn diagrams in which the area of each region is in proportion to the size of the population that it represents are said to be *area proportional*, and this is now an active area of research; see, for example, [11] and [21]. It is also of interest to produce diagrams in which each region has the same area, particularly if the shape of each region is identical. Such diagrams might be useful because of their visual properties, or because of the ease of labelling them. For example, the article [7] in a microbiology journal used a minimal area Venn diagram from [11] for data visualization.

These diagrams use the perimeters of polyominoes as the curves of the Venn diagram.

Definition 6.0.5. A *polyomino*, "invented" by Solomon Golomb [14], is a set of unit squares whose corners lie on the integer lattice $\mathbf{Z}^2$ and whose perimeter forms a single simple closed curve.

Perhaps the first such diagrams appeared on the website of Mark Thompson [22], which introduced the idea of using the perimeters of polyominoes as the curves of Venn diagrams, with the additional property that each region is a unit square. Chow and Ruskey[11] further investigated the question of minimizing the number of unit squares, and asked, but did not resolve, the question of whether polyVenn diagrams exist for any number of curves.

In this chapter, we settle several conjectures/questions from [11], by showing that Venn diagrams exist where each region is represented by a unit square. Most notably we show how to construct minimum area polyomino diagrams in bounding rectangles

of size $2^r \times 2^c$ whenever $r, c \geq 2$. We also show that we can use the half-set system from Chapter 5 to demonstrate two different expansion techniques: one where a 4×2^c minimum area Venn diagram is expanded to a $4 \times 2^{c+3}$ diagram, the other where a $2^r \times 2^c$ minimum area Venn diagram is expanded to a $2^{r+r'} \times 2^{c+c'}$ diagram.

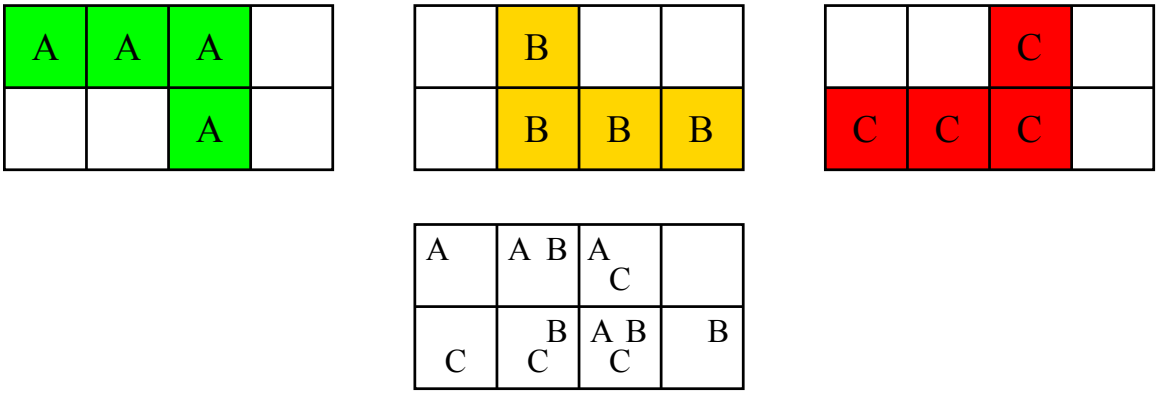

Figure 6.1: A $(1, 2)$-polyVenn

6.1 Definitions

Consider the set of three L shaped tetrominoes A, B, C, shown at the top of Figure 6.1. If these are overlapped on top of each other, then all subsets of {A,B,C} occur as a unique unit square in the result, the 2×4 rectangle shown at the bottom of Figure 6.1. In other words, the curves that comprise the perimeters of these tetrominoes form a Venn diagram when overlaid as shown.

Although Golomb allows for it, for our purposes, a polyomino does not contain any holes, i.e, they are *simply connected*. We assume that unit squares, and therefore the polyominoes themselves, include their interiors.

Definition 6.1.1. A *polyomino Venn diagram*, or *polyVenn* consists of polyominoes $P_1, \ldots, P_n$, such that each of the 2^n sets $X_1 \cap X_2 \cap \ldots \cap X_n$ is a nonempty and simply connected region, where X_i is either part of P_i or not part of P_i.

Note that for any set of polyominoes that yield a polyVenn, its perimeters make up the Jordan curves of a Venn diagram. In this thesis, we require that all Jordan curves intersect each other at a finite number of points. Clearly, the boundaries of the polyominoes are not Venn diagrams in the classical definition of Grunbaum and this thesis, but the adjustment is clear.

Chow and Ruskey [11] introduced the concept of a *minimum* area n-polyVenn, where the total area of the n intersecting polyominoes is $2^n - 1$ unit squares. We modify their definition by focusing on the polyomino as a simply connected subset of a larger set $\mathbb{G}$ of unit squares in $\mathbb{R}^2$. The interior of a polyomino, P, is comprised of all its unit squares while any squares in $\overline{P} = \mathbb{G} \setminus P$ form the exterior of P. Any portion of $\mathbb{R}^2$ that is not in $\mathbb{G}$ is ignored.

Definition 6.1.2. A *minimum area polyomino Venn diagram* is a polyVenn diagram, together with a set of unit squares $\mathbb{G}$, where each of the regions defined in Definition 6.1.1 is a single unit square in $\mathbb{G}$.

The set $\mathbb{G}$ provides a consistency in that the empty set is represented by a bounded region in the same way as the other regions. We are primarily concerned with polyVenns where $\mathbb{G}$ is a bounding rectangle.

Definition 6.1.3. An (r, c)-polyVenn is an n-set minimum area polyVenn, with $n = r + c$ and $r, c > 0$, in which $\mathbb{G}$ is a $2^r \times 2^c$ rectangle.

By convention, we will only consider (r, c)-polyVenns with $r \leq c$, since if this is not the case, then the diagram can be rotated 90° to make it so. Here further, we refer to minimum area polyomino Venn diagrams as *polyVenns*, *polyVenn diagrams* or *n-set polyVenns*.

6.1.1 PolyVenns and HSSs

The polyVenn in Figure 6.1 is an example of a $(1, 2)$-polyVenn. Note that in an n-set polyVenn, each polyomino contains exactly 2^{n-1} unit squares, and its intersection with every other polyomino is exactly 2^{n-2} squares.

Theorem 6.1.4. *An (r, c)-polyVenn is an HSS on a set of unit squares.*

Proof. Let $\mathbf{P}$ be an (r, c)-polyVenn. Then $\mathbb{G} = \{g_{ij} \mid 1 \leq i \leq 2^r, 1 \leq i \leq 2^c\}$. Also, $\mathbf{P}$ is a set of polyominoes $\{P_1, P_2, \ldots, P_n\}$ where each $P_i \subseteq \mathbb{G}$. By Definition 6.1.2, each of the regions defined in Definition 6.1.1 is a single unit square. That means that for any set $B \subseteq \{1, 2, \ldots 2^{r+c}\}$,

$$\left(\bigcap_{i \in B} P_i\right) \cap \left(\bigcap_{i \notin B} \overline{P_i}\right) = g,$$

where $g \in \mathbb{G}$. Therefore, by Theorem 5.0.2, $\mathbf{P}$ is an HSS. □

Naturally, for any polyomino P_i on a polyVenn, its set of unit squares must form a simply connected unit in order for the perimeter of P_i to be a curve on the Venn diagram. If $\overline{P_i} = \mathbb{G} \setminus P_i$ is also a set of simply connected unit squares, then we say that P_i is *self-complementing*.

Corollary 6.1.5. *If an (r, c)-polyVenn $\mathbf{P}$ contains a self-complementing polyomino, P_i, then the set $\overline{P_i}$ can be substituted for P_i and the result is an (r, c)-polyVenn.*

Proof. This follows from the proof of Lemma 5.1.1. □

6.2 Expanding an existing diagram

It is natural to ask whether (r, c)-polyVenns exist for all r and c. If they do, is there a method to construct them? We want to know if it is possible to extend a polyVenn by stretching the diagram and adding more polyomino pieces.

Our primary tool for approaching this question is the *expansion* of the HSS, shown in Lemma 5.0.3. The idea is to take an existing (r, c)-polyVenn $\mathbf{P}$ and expand it into an $(r + r', c + c')$-polyVenn by the addition of $r' + c'$ polyominoes. Both $\mathbb{G}$ and the set of polyominoes are "expanded" by uniformly stretching vertically by a factor of $2^{r'}$ and horizontally by a factor of $2^{c'}$. The key element is finding the $r' + c'$ new polyominoes and Lemma 5.0.3 provides us with a method to do this.

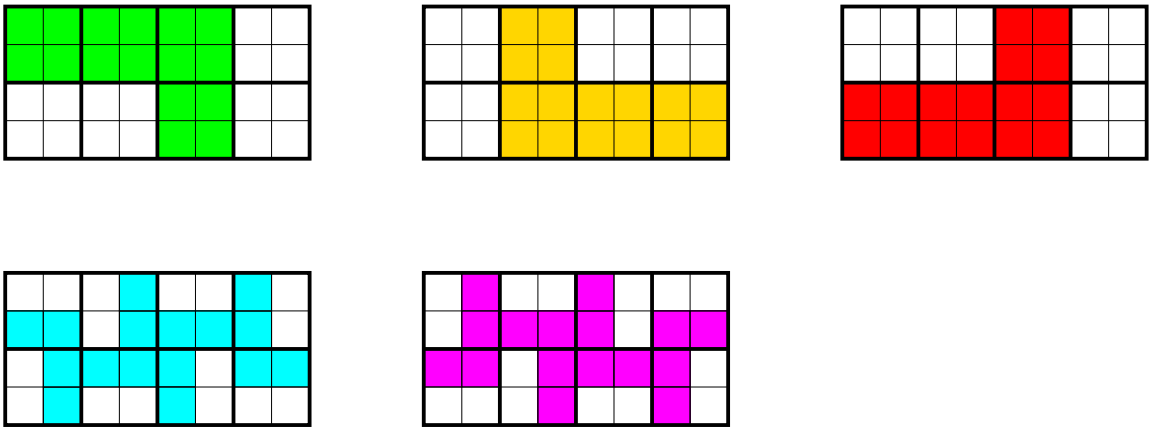

Figure 6.2: A $(2, 3)$-polyVenn created by expansion.

Figure 6.2 shows an expansion of the $(1, 2)$-polyVenn of Figure 6.1 into a $(2, 3)$-polyVenn. Each polyomino is positioned on the 4×8 expanded rectangular grid, the original grid being $\mathbb{G} = \{g_{ij} \mid 1 \leq i \leq 2, 1 \leq j \leq 4\}$. The expanded grid $X(\mathbb{G})$ is

comprised of eight *minigrids*, which in this case are the 2×2 expansions of the original set of unit squares of $\mathbb{G}$. The top three polyominoes are the expanded originals from Figure 6.1. The bottom two polyominoes have the property that they contain exactly one domino for each minigrid in the diagram. The two dominoes within each minigrid make up a single $(1,1)$-polyVenn.

Although the example shows it, the creation of the $r' + c'$ new polyominoes does not require that the intersection of each minigrid with the new polyominoes forms an (r', c')-polyVenn. The polyominoes must be simply connected on the larger grid $X(\mathbb{G})$, but the collections of squares within a minigrid can be disconnected, in other words, the set of subsets on each minigrid must be an HSS.

Theorem 6.2.1. *Let* $\mathbf{P}$ *be a polyVenn on a* $2^r \times 2^c$ *grid* $\mathbb{G}$*. Let* $\mathbb{G}$ *be expanded to* $X(\mathbb{G})$ *where every grid* $g_{ij} \in \mathbb{G}$ *is replaced by a* $2^{r'} \times 2^{c'}$ *expansion,* G_{ij}*,* $0 \leq i < 2^r$*,* $0 \leq j < 2^c$*. Suppose, for every* G_{ij}*, there is an HSS,* $\mathbf{M}_{ij} = (M_{ij1}, M_{ij2}, \ldots, M_{ij(r'+c')})$ *on the unit squares of* G_{ij} *for which the following is true:*

For each $1 \leq k \leq r' + c'$*,*

$$E_k = \bigcup_{\substack{0 \leq i < 2^r \\ 0 \leq j < 2^c}} M_{ijk} \quad \textit{covers a simply connected space on } X(\mathbb{G}).$$

Then an $(r + r', c + c')$*-polyVenn exists that is a* $2^{r'} \times 2^{c'}$ *expansion of* $\mathbf{P}$*.*

Proof. We use i, j and k as defined in the theorem statement. Let q be defined as an integer such that $1 \leq q \leq r + c$. We know, by Theorem 6.1.4 that $\mathbf{P} = \{P_1, P_2, \ldots, P_{r+c}\}$ is an HSS on $\mathbb{G}$. The expansion of these polyominoes is $X(\mathbf{P}) = \{X(P_1), X(P_2), \ldots, X(P_{r+c})\}$ where every g_{ij} in P_q corresponds to G_{ij} in $X(P_q)$. Since G_{ij} is a rectangle for all and P_q is also simply connected, the set of unit squares $X(P_q)$ is also a polyomino. By Lemma 5.0.3, we know that $X(\mathbf{P}) \cup E$ is an HSS on $X(\mathbb{G})$. This means that for all q and k, the intersection of P_q or $\overline{P_q}$ and E_k or $\overline{E_k}$ is a single unit square. Since E_k is defined as a simply connected set by hypothesis, each of these are polyominoes on $X(\mathbb{G})$ and we have a $(r + r', c + c')$-polyVenn. □

6.3 Two expansions

We have two expansion results for (r, c)-polyVenns, Theorem 6.3.1 and Theorem 6.3.2.

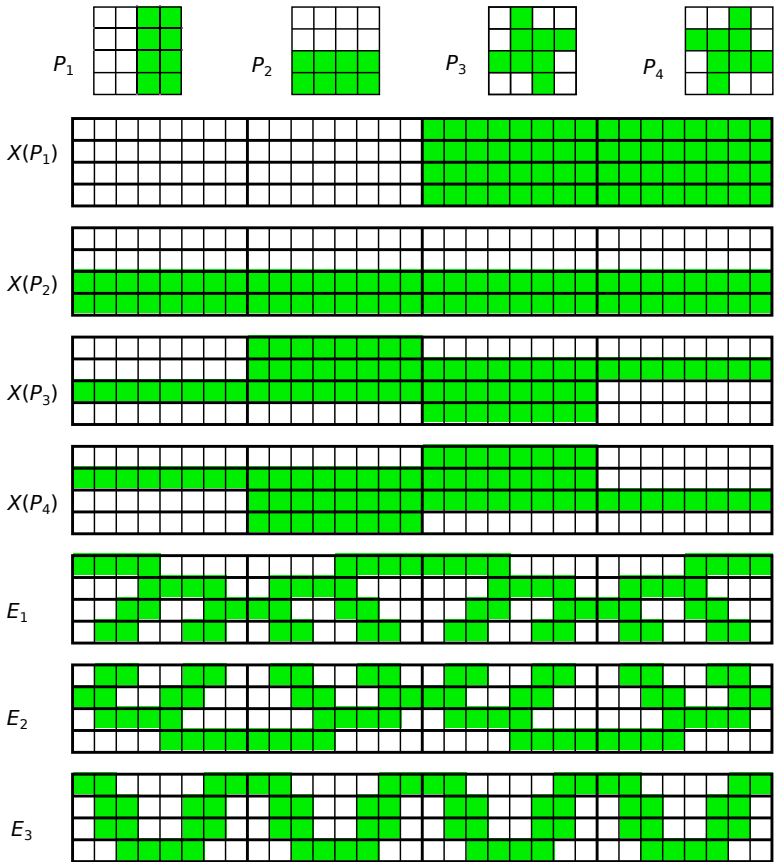

Figure 6.3: A 4×8 expansion of a $(2,2)$-polyVenn: the top 3 pieces are $\{E_1, E_2, E_3\}$; the bottom pieces are $\{X(P_1), X(P_2), X(P_3), X(P_4)\}$.

Theorem 6.3.1. *For $c \geq 0$, a $(2,c)$-polyVenn can be expanded to a $(2,c{+}3)$-polyVenn.*

Proof. Let $\mathbf{P}$ be a $(2,c)$-polyVenn. We will exhibit a 1×8 expansion by constructing sets E_1, E_2 and E_3 such that the conditions of Theorem 6.2.1 are satisfied. It is convenient to describe the following 8 sets of subsets of the expanded unit squares of $\mathbf{V}$ as drawings.

For each g_{ij} in $\mathbb{G}$, let G_{ij} be a 1×8 mini grid on $X(\mathbb{G})$, the $2^r \times 2^{c+3}$ expansion of $\mathbb{G}$. Consider a set $\mathbf{M}_{ij} = \{M_{ij1}, M_{ij2}, M_{ij3}\}$, defined by one of the following eight triples of a 1×8 minigrid:

$\mathbf{M}_{11} =$ (, ,), $\mathbf{M}_{12} =$ (, ,),
$\mathbf{M}_{21} =$ (, ,), $\mathbf{M}_{22} =$ (, ,),
$\mathbf{M}_{31} =$ (, ,), $\mathbf{M}_{32} =$ (, ,),
$\mathbf{M}_{41} =$ (, ,), $\mathbf{M}_{42} =$ (, ,).

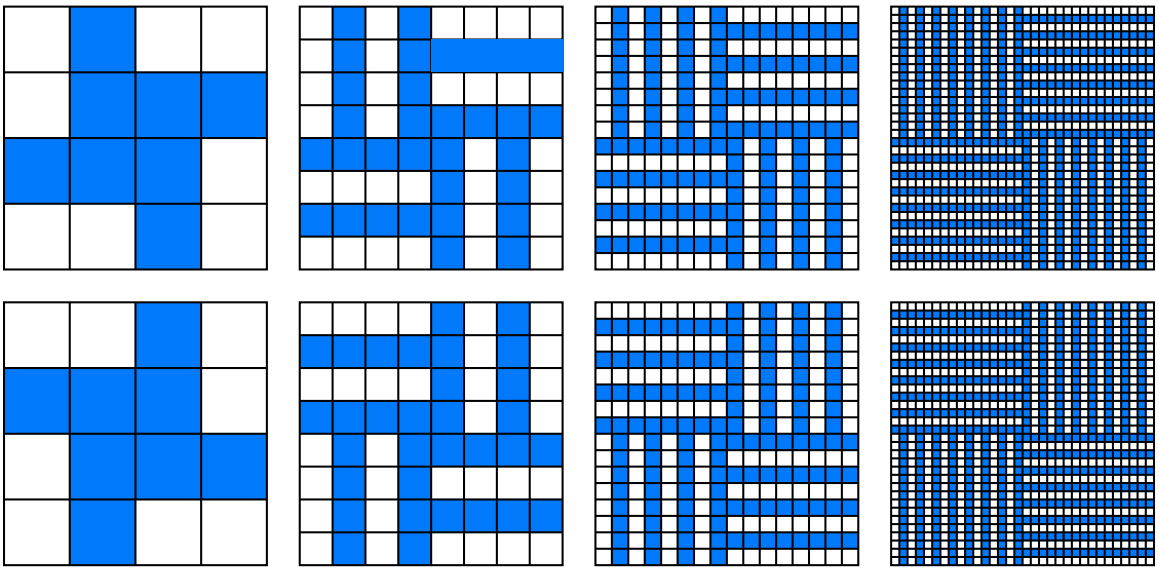

Figure 6.4: The constructed E_1 and E_2 described in the proof of Theorem 6.3.2, shown for $(r, c) = (1,1),(2,2),(3,3)$ and $(4,4)$. Note that the construction is valid for any $r, c \geq 1$.

We remark that each component of each $\mathbf{M}_{i1}$ is a horizontal reflection of the same component of $\mathbf{M}_{i2}$. It is easy to check from the little grid drawings that each $\mathbf{M}_{ij}$ is an HSS. To create the new polyominoes, place $\mathbf{M}_{i1k}$ into the minigrid G_{ij} of E_k if j is odd or place $\mathbf{M}_{i2k}$ if j is even.

It remains to be shown that E_1, E_2 and E_3 are simply connected. This can be verified by inspection for $c = 0, 1$. For larger c, the resulting polyominoes are formed by repeatedly concatenating the $c = 1$ case polyominoes with themselves. Again by inspection, this pattern produces simply-connected shapes as illustrated in Figure 6.3. □

Theorem 6.3.2. *For $r, c \geq 1$, an (r, c)-polyVenn can be expanded to an $(r+1, c+1)$-polyVenn.*

Proof. Let $\mathbf{P}$ be an (r, c)-polyVenn on a $2^r \times 2^c$ grid $\mathbb{G}$. We will exhibit a 2×2 expansion by constructing sets E_1 and E_2 such that the conditions of Theorem 6.2.1 are satisfied. Similar to the previous proof, we describe sets of subsets using minigrid drawings.

For each g_{ij} in $\mathbb{G}$, let G_{ij} be a 2×2 minigrid on $X(\mathbb{G})$, the $2^{r+1} \times 2^{c+1}$ expansion of $\mathbb{G}$. Consider a set $\mathbf{M}_{ij} = \{M_{ij1}, M_{ij2}\}$ defined by one of the following four sets of

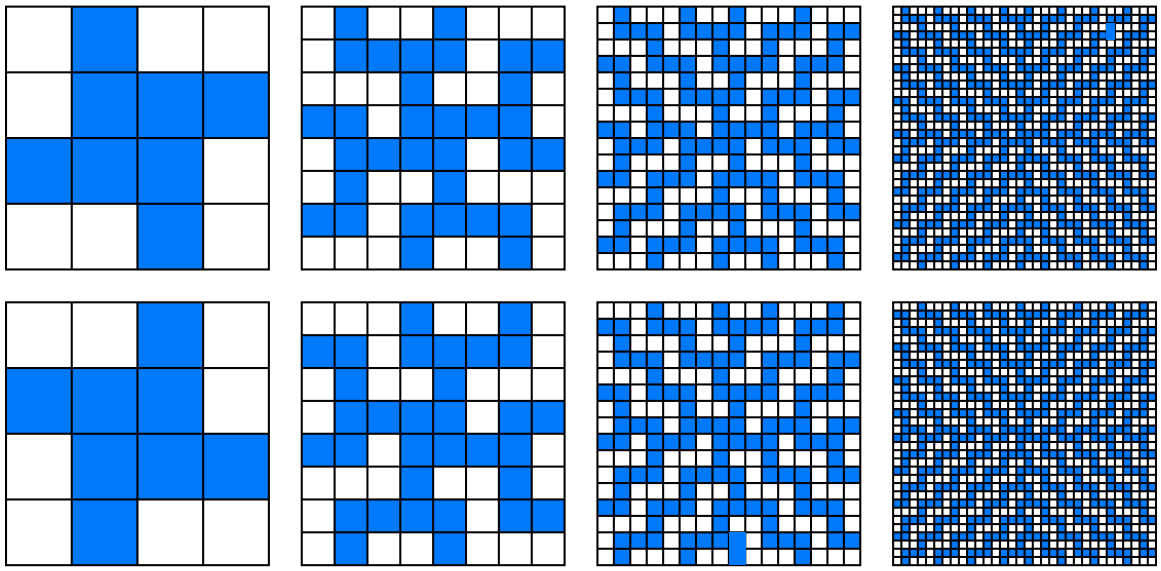

Figure 6.5: Alternative E_1 and E_2 polyominoes for (r, c) = (1,1),(2,2),(3,3) and (4,4). This construction is also valid for any $r, c \geq 1$.

a 2×2 minigrid:

$$\mathbf{M}_{nw} = \{ \blacksquare , \blacksquare \} \qquad \mathbf{M}_{ne} = \{ \blacksquare , \blacksquare \}$$
$$\mathbf{M}_{sw} = \{ \blacksquare , \blacksquare \} \qquad \mathbf{M}_{se} = \{ \blacksquare , \blacksquare \}$$

It is easy to see that each set of subsets is a $(1,1)$-polyVenn, and thus an HSS.

To create the pieces E_1 and E_2, for the minigrid G_{ij}, set (ij) to be $(nw), (ne), (sw)$ or (se), depending on whether G_{ij} is in the upper left, upper right, lower left or lower right quadrant of $X(\mathbb{G})$. We provide two examples of placement, illustrated in Figures 6.4 and 6.5. Figure 6.4 always adds $\mathbf{M}_{ij1}$, to E_1 and $\mathbf{M}_{ij2}$ to E_2; Figure 6.5 adds the vertical subset of $\mathbf{M}_{ij}$ to E_k at minigrid G_{ij} if $k + i + j$ is odd and the horizontal subset if $k + i + j$ is even.

We prove that E_1 and E_2, illustrated in Figure 6.4, are each simply connected and leave it to the reader to see that a similar proof holds for Figure 6.5. First we show that E_k is connected for $k = \{1, 2\}$. Note that E_1 and E_2 are both horizontal and vertical reflections of each other and on E_k, every quadrant is rotationally identical, so whatever is true for one quadrant is true for all four. Consider a unit square $u \in E_k \cap G_{ij}$ such that $1 \leq i \leq 2^{r-1}$ and $1 \leq j \leq 2^{c-1}$, in other words u is in the left upper quadrant of E_k. Let $v \neq u \in M_{ijk}$, be another unit square in the left upper quadrant of E_k. Then v is either vertically or horizontally aligned with u on G_{ij}.

Suppose $\{u, v\}$ is a vertically aligned subset. Then $k = 1$ because $M_{nw1} =$. On E_1, we also know that $M_{se1} =$. Therefore there is a vertical path from u to a unit square in $G_{i,2^{r-1}}$, which is then connected to the centre of E_1. A similar argument applies if $\{u, v\}$ is horizontally aligned. If u is connected to the centre from any quadrant of E_1 or E_2 which can be reflected or rotated to the left upper quadrant of E_1, then E_k is connected.

To prove that E_k is simply connected we show that $\overline{E_k}$ is also connected. Consider a unit square w that is not in E_k. Using a similar argument, assume $w \in E_1$ and $w \in M_{nw}$. Then $w \in$ which is the complement of , the subset for all $G_{ij} \cap E_1$ and there is a vertical path from w to the exterior of E_1. Thus $\overline{E_k}$ is connected and E_k is simply connected. □

Figure 6.5 is similarly proven to be simply-connected, except that the paths between u and the centre point and between w and the outer grid border are zigzagged.

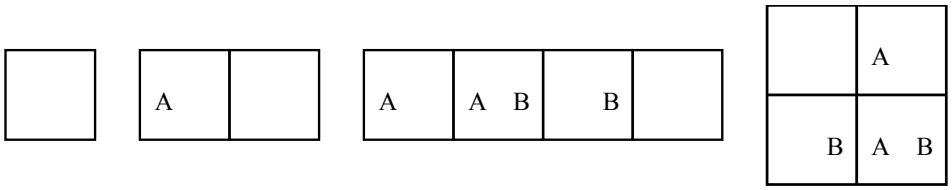

Figure 6.6: The trivial polyVenns.

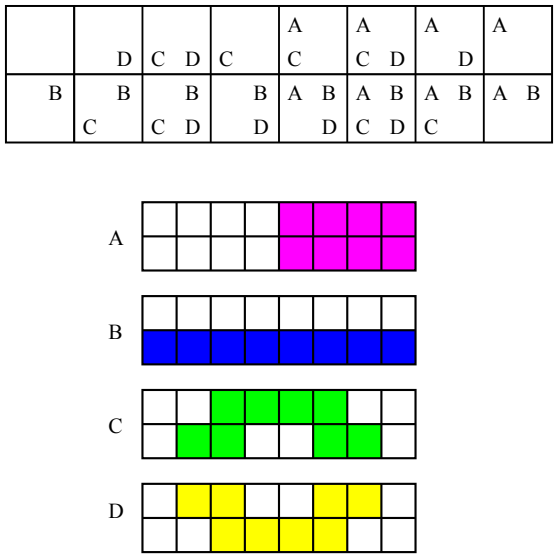

Figure 6.7: A $(1,3)$-polyVenn with nice symmetry

6.4 The base cases

Theorems 6.3.1 and 6.3.2 provide a way of inductively determining ranges of (r,c) pairs for which an (r,c)-polyVenn exists. But in order to begin we must establish some base cases.

It is sufficient to consider base cases with $r = 0$ and $r = 1$. Figure 6.6 illustrates the trivial cases for the 1×1 *empty* polyVenn, $c = 0, 1, 2$ and the $(1,1)$-polyVenn. The trivial $(1,2)$-polyVenn was shown in Figure 6.1.

Proposition 6.4.1. *A $(0,c)$-polyVenn does not exist for any $c > 2$.*

Proof. Consider a $(0,c)$-polyVenn where $c > 2$. Each of the c polyomino pieces consists of 2^{c-1} consecutive squares. Without loss of generality, let P_1 be the leftmost polyomino. All other polyominoes must each intersect P_1 in its rightmost 2^{c-2} squares. There is only one polyomino consisting of 2^{c-1} horizontally consecutive squares that does this. A polyVenn cannot have two totally overlapping polyominoes, so we must have $c \leq 2$. □

D	C D	BC	BC D	ABC	ABC E	ABC D	A C D	ABC D E	A C D E	C D E	BC E	C E	A C E	A C	C
B	B D	AB D	A D	A D E	D E	B D E	B E	BC D E	AB D E	AB E	AB	A	A E	E	

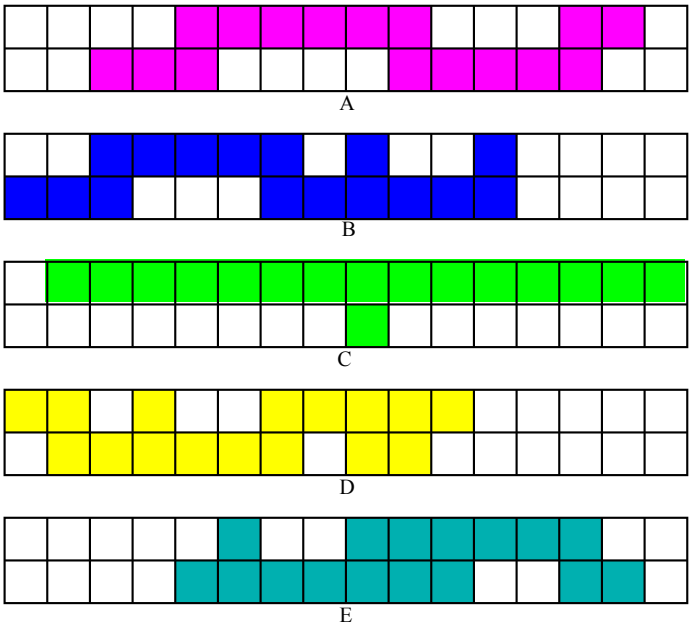

Figure 6.8: A $(1, 4)$-polyVenn

In Figures 6.7 and 6.8, we show a $(1,3)$-polyVenn and a $(1,4)$-polyVenn, respectively. For $c > 4$ we do not know of any $(1,c)$-polyVenn and we make the following conjecture.

Conjecture 6.4.2. *A $(1,c)$-polyVenn does not exist when $c > 4$.*

Our motivation for this conjecture comes from the observation that having only two rows is very constraining in the construction of polyVenns. For example, no $(0,1)$ expansion seems possible on an existing $(1,c)$-polyVenn. A partial computer search has not produced a $(1,5)$-polyVenn. The problem size is currently too large to determine conclusively that Conjecture 6.4.2 is true.

If $\text{num}(n)$ is the number of polyomino pieces with area 2^{n-1} and height no more than two, up to vertical and horizontal reflection, then $\text{num}(2) = 2$, $\text{num}(3) = 6$, $\text{num}(4) = 63$, $\text{num}(5) = 8189$ and $\text{num}(6) = 140{,}473{,}849$. We note that the polyominoes themselves become somewhat restricted in their shapes as the following lemma shows.

Lemma 6.4.3. *There is never a rectangular polyomino that is part of a $(1,c)$-polyVenn for $c > 3$. A rectangular polyomino can exist when $c = 3$ (see Figure 6.9), but its edge must be within one column of the border of its base region.*

Proof. Let $\mathbf{P}$ be a $(1,c)$-polyVenn containing the rectangular polyomino piece P_1. Then P_1 contains a subset of the unit squares $g_{ij} \in \mathbb{G}$, where $1 \leq i \leq 2$ and $1 \leq j \leq 2^{n-1}$. Because of the horizontal symmetry, we assume that the leftmost unit squares of P_1 are g_{1s} and g_{2s} where $s \leq 2^{n-3}+1$. Because of its position, $\overline{P_1}$ contains r squares to the right of P_1, where $2^{n-3} \leq r \leq 2^{n-2}$. We refer to these squares in $\overline{P_1}$ as the *right complement squares*, while the squares to the left of P_1, if they exist, are called the left complement squares.

We organize this proof as a series of facts. P_1 will always refer to the rectangular piece and P_i, P_j, P_k are some polyominoes in $\mathbf{P}$, where $i \neq j \neq k$ and $2 \leq i, j, k \leq n$.

Fact 1: When $s \leq 2^{n-3}$, then a polyomino P_i contains some of the right complement squares.

- This is because the number of left complement squares is less than 2^{n-2} and $|\overline{P_1} \cap P_i| = 2^{n-2}$.

Fact 2: Only one polyomino P_i can span all the columns of $\mathbb{G}$ that are numbered $s, s_1, \ldots, s_2^{n-2} - 1$.

- Since $|P_1 \cap P_i| = 2^{n-2}$ and P_1 spans 2^{n-2} columns, P_i will contain either the top or bottom rows of columns $s, s+1, \ldots, s+2^{n-2}-1$. If another polyomino P_j spans the columns then $P_1 \cap P_i \cap P_j$ is either P_1 or $\emptyset$ and $\mathbf{V}$ cannot be a polyVenn.

Fact 3: If a polyomino P_i contains only left complement squares, then $n = 3$.

- Let P_j, be a third polyomino in $\mathbf{P}$. If P_j contains only left or right complement squares, then $|\overline{P_1} \cap P_i \cap P_j| = 2^{n-2}$ or 0, respectively. So, P_j must span all the columns of P_1 and by Fact 2, P_j it is the only polyomino to do so.

Fact 4: When $n > 3$, every polyomino other than P_1 must contain right complement squares and there are only zero or two left complement squares.

- The first part of the claim is clear from Fact 3. If the number of left complement squares is greater than zero, then it is an even number and only one square can be the "empty set" square, so at least one other square must belong to polyomino P_i. By Fact 2, P_i is the only polyomino that can contain left complement squares and likewise, for any P_k, $P_i \cap \overline{P_1} \cap \overline{P_k}$ must also be a single square. Thus, there are 2 left complement squares. Figure 6.9 shows examples of this fact.

Now suppose $n > 4$: Since by Facts 2 and 4, only one polyomino P_i beside P_1 can contain a square in column s; without loss of generality, let it be the top square, g_{1s}. Then g_{1s} and $g_{1(s+1)}$ are both in P_i. Since all the squares of column s and $s+1$ are in P_1, by Definition 6.1.2, either a single polyomino P_j contains both squares in column $s+1$ or at least 2 polyominoes, P_j and P_k contain squares from column $s+1$. It is easy to see that the first case cannot happen: If P_j contains both squares of column $s+1$ then it must span P_1 for the rest of the columns so that $P_1 \cap P_j = 2^{n-2}$. If it spans the top squares then $|P_1 \cap P_i \cap P_j| = 2^{n-2}$; if it spans the bottom squares then $|P_1 \cap P_i \cap P_j| = 1$ and $n = 3$. Therefore, the squares of column $s+1$ are contained in at least 2 polyominoes, P_j and P_k.

By Fact 4, both P_j and P_k must span $2^{n-2}-1$ rows of P_1. Because $|P_1 \cap P_j| = 2^{n-2}$, P_j must contain the upper and lower squares of exactly one column. More specifically, since $|P_1 \cap P_i \cap P_j| = 2^{n-3}$, P_j contains both squares in column $s+2^{n-3}$. The same requirements hold for P_k. Thus P_j and P_k are vertical reflections of each other and

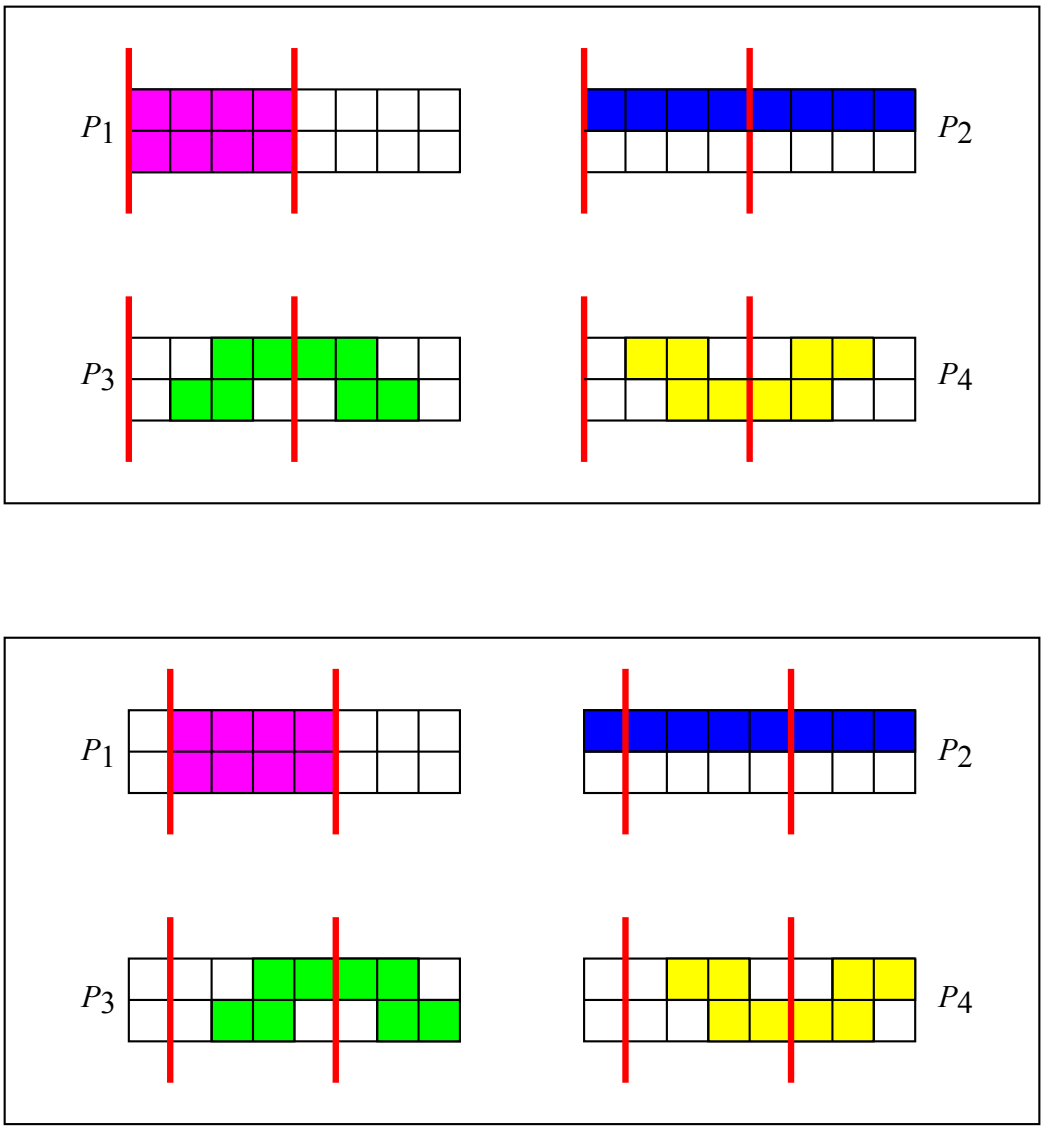

Figure 6.9: Two placements for a rectangular polyomino in a $(1, 3)$-polyVenn

$g_{1,(s+2^{n-3})} = P_1 \cap P_i \cap P_j \cap P_k$. Adding a 5^{th} polyomino, P_5 would require that either $g_{1,(s+2^{n-3})} \in P_5$ and then $P_1 \cap P_2 \cap P_3 \cap P_4 \cap \overline{P_5} = \emptyset$, or $g_{1,(s+2^{n-3})} \notin P_5$ and $P_1 \cap P_2 \cap P_3 \cap P_4 \cap P_5 = \emptyset$. Therefore, for $n = 1 + c > 4$, no rectangular polyomino exists in $\mathbf{V}$. □

6.5 Summary of results

Given the base cases and the expansion theorems, we summarize the known polyVenns in the following theorem:

Theorem 6.5.1. *For every $n > 0$ and for every $r, c > 2$ where $r + c = n$, there exists a minimum polyomino Venn diagram bounded in a $2^r \times 2^c$ bounding rectangle.*

Proof. Applying Theorems 6.3.1 and 6.3.2 with the bases cases in Figure 6.6 completes the proof. □

Table 6.1 shows how the base cases are extended and summarizes our state of knowledge about for which r and c an (r, c)-polyVenn exists.

Table 6.1: Our knowledge of existence of (r, c)-polyVenns. Rows and columns are labelled with bounding rectangle dimensions (2^r and 2^c). Fn indicates existence by example given in Figure n; '×' indicates nonexistence following from Proposition 6.4.1; A and B are configurations implied by Theorems 6.3.1 and 6.3.2, respectively; Question marks are currently open cases.

	2^0	2^1	2^2	2^3	2^4	2^5	2^6	2^7	2^8	$\cdots$
2^0	F6.6	F6.6	F6.6	×	×	×	×	×	×	$\cdots$
2^1		F6.6	F6.6	F6.7	F6.8	?	?	?	?	$\cdots$
2^2			B	B	B	A	A	A	A	$\cdots$
2^3				B	B	B	B	B	B	$\cdots$
2^4					B	B	B	B	B	$\cdots$
2^5						B	B	B	B	$\cdots$
2^6							B	B	B	$\cdots$
2^7								B	B	$\cdots$
2^8									B	$\cdots$
$\vdots$										$\ddots$

6.6 Rectangles that omit the empty set

In this section we discuss the problem of finding n-set polyVenns in which each polyomino fits into into an $h \times w$ rectangle where $2^n - 1 = hw$, so that the region for the empty set (the intersection of the exterior of all the polyomiones) is the exterior of the rectangle. We refer to such polyVenns as minimum area empty-set-omitting polyVenns or *ESO-polyVenns*. For $n > 2$, it is only possible for such polyVenns to exist when $M_n = 2^n - 1$ is *not* a Mersenne prime, since otherwise the resulting diagram must have dimensions $1 \times M_n$ and would imply the existence of a polyVenn that violates Proposition 6.4.1. For $n > 2$, the first eight such numbers are $2^4 - 1 = 15 = 3 \cdot 5$,

$2^6-1 = 63 = 3^2 \cdot 7$, $2^8-1 = 255 = 3 \cdot 5 \cdot 17$, $2^9-1 = 511 = 7 \cdot 73$, $2^{10}-1 = 1023 = 3 \cdot 11 \cdot 31$, $2^{11}-1 = 2047 = 23 \cdot 89$, $2^{12}-1 = 4095 = 3^2 \cdot 5 \cdot 7 \cdot 13$. This problem was first mentioned in [11].

Theorem 6.6.1. *For all $r \geq 0$, there exists a minimum area ESO-polyVenn with a $(2^r - 1) \times (2^r + 1)$ base rectangle.*

Proof. We use a simple construction of a polyVenn of dimension $(2^r - 1) \times (2^r + 1)$, where $r = n/2$, based on the successive 2×2 expansions of the $(1, 1)$-polyVenn of Figure 6.6. We can use either construction detailed in the proof of Theorem 6.3.2 and illustrated in Figures 6.4 and 6.5. In either case, the following actions are easily verified as possible on the the resulting (r, r)-polyVenns: First, the empty set, represented by the upper left square of the bounding rectangle, is removed from the diagram. Second, the top row, minus the empty set square, is removed, rotated 90° clockwise, and attached to the right side of the diagram.

The resulting arranged polyominoes continue to be simply-connected and thus polyominoes. Performing the rearrangement described results in a minimum area polyVenn whose bounding rectangle is a $(2^r - 1) \times (2^r + 1)$ with the empty set removed in the first step. □

Figure 6.10: Converting the base case $(1, 1)$-polyVenn to fit onto a 1×3 bounding rectangle, omitting the empty set.

Figure 6.10 demonstrates how this construction converts the $(1, 1)$-polyVenn to a polyVenn on a 1×3 base rectangle. Figure 6.11 shows the result of the construction when $n = 8$ and the $(4, 4)$-polyVenn was constructed using repeated 2×2 expansions of the type illustrated in Figure 6.5.

We do not know of any minimum area ESO-polyVenns base rectangle dimensions not covered by this construction. The next few such dimensions are

$$63 = 3 \times 21,\ 255 = 3 \times 85 = 5 \times 51, \text{ and } 511 = 7 \times 73.$$

Given the seemingly mysterious nature of the Mersenne primes, we feel that there will not be any general construction possible when n is odd and M_n is not prime.

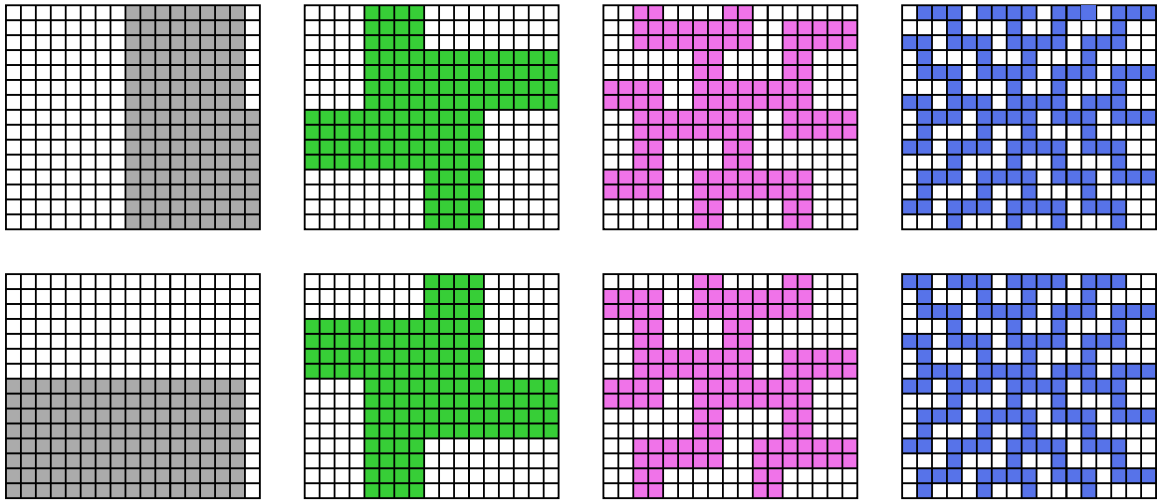

Figure 6.11: Omiting the empty set: a polyVenn on the $(2^{n/2}-1) \times (2^{n/2}+1)$ grid for $n = 8$.

Chapter 7

Future Research

In this thesis, we settled some open problems and made some progress on others. Winkler's conjecture remains open for simple, irreducible Venn diagrams. We have shown that it is true for all simple Venn diagrams for up to and including five curves. It is also true for a subset of irreducible Venn diagrams for all values of $n > 5$. Although it does not apply directly to polyVenn diagrams, we can state the following modification, which was proven in Chapter 6.

Fact 7.0.2. *Every minimum area polyomino Venn diagram with n polyominoes is extendible to a $n + 2$ minimum area polyomino Venn diagram.*

In our opinion, it is unlikely that we can extend a polyVenn by the addition of a single polyomino. However, this raises the question of whether one can decide whether an independent set of n curves can be extended to a Venn diagram of $n + k$ curves, where k is the minimum number of curves needed to create a Venn diagram.

Several conjectures, 3.1.5, 3.2.1 and 3.2.2, were made about face-balanced diagrams and it would be nice to be able to answer them. We believe it would be beneficial to look at diagrams with regions exceeding 32. A computer search that generates them directly is needed; generating all quadangulations with more than 32 vertices and then filtering them is far too inefficient. We know that face-balanced diagrams are a superset of Venn diagrams, but we are still unsure of just how closely related they are. One question we have not answered is whether all their dual graphs have Hamilton cycles. At the time of the writing of this concluding chapter, we conjecture that they do not.

Bibliography

[1] Agarwal, P. K. and M. Sharir. "Pseudo-line arrangements: Duality, Algorithms, and Applications". *SIAM Journal on Computing*, 34 (2005) 526–552.

[2] Brinkmann, G. and B. D. McKay. "Fast generation of planar graphs". *Match-Communications in Mathematical and in Computer Chemistry*, 58 (2007) 323–357.

[3] Brinkmann, G. and B. D. McKay. "plantri" (2011). URL `http://cs.anu.edu.au/~bdm/plantri/`. Software to generate planar graphs.

[4] Bultena, B., B. Grünbaum, and F. Ruskey. "Convex Drawings of Intersecting Families of Simple Closed Curves". *11th Canadian Conference of Computational Geometry*, (1999) 18–21.

[5] Bultena, B., M. Klimesh, and F. Ruskey. "Minimum Area Polyomino Venn Diagrams". *Journal of Computational Geometry*, 3 (2012), no. 1 154–167.

[6] Bultena, B. and F. Ruskey. "Venn Diagrams with Few Vertices". *Electronic Journal of Combinatorics*, 5 (1998) R44.

[7] Casimiro, S., R. Tenreiro, and A. Monteiro. "Identification of pathogenesis-related ESTs in the crucifer downy mildew oomycete Hyaloperonospora parasitica by high-throughput differential display of distinct phonotpic interactions with Brassica oleracea". *Journal of Microbiological Methods*, 66 (2006) 466–478.

[8] Chilakamarri, K. B., P. Hamburger, and R. E. Pippert. "Hamilton Cycles in Planar Graphs and Venn Diagrams". *Journal of Combinatorial Theory*, 67 (1996), no. 0047 296–303.

[9] Chilakamarri, K. B., P. Hamburger, and R. E. Pippert. "Venn diagrams and planar graphs". *Geometriae Dedicata*, 62 (1996) 73–91.

[10] Chilakamarri, K. B., P. Hamburger, and R. E. Pippert. "Analysis of Venn Diagrams Using Cycles in Graphs". *Geometriae Dedicata*, 82 (2000) 193–223.

[11] Chow, S. and F. Ruskey. "Minimum Area Venn Diagrams Whose Curves are Polyominoes". *Mathematics Magazine*, 80 (2007), no. 2 91–103.

[12] Edwards, A. W. "Seven-Set Venn Diagrams with Rotational and Polar Symmetry". *Combinatorics, Probability and Computing*, 7 (1998) 149–152.

[13] Garey, M. R. and D. S. Johnson. *Computers and intractability: a guide to the theory of NP-completeness*. San Francisco: W. H. Freeman and Company (1979).

[14] Golomb, S. W. *Polyominoes*. Princeton, New Jersey: Princeton Science Library (1994).

[15] Grünbaum, B. "The Construction of Venn Diagrams". *The College Mathematics Journal*, 15 (1984), no. 3 238–247.

[16] Grünbaum, B. "Venn Diagrams I". *Geominatorics*, 1 (1992) 5–12.

[17] Hales, T. C. "The Jordan Theorem, Formally and Informally". *The American Mathematical Monthly*, 114 (2007), no. 10 882–894.

[18] Hamburger, P. and R. E. Pippert. "Simple, Reducible Venn Diagrams on Five Curves and Hamiltonian Cycles". *Geometriae Dedicata*, 68 (1997) 245–262.

[19] Mamakani, K. and F. Ruskey. "A New Rose: The First Simple Symmetric 11-Venn Diagram". *arXiv.org 1207.6452*, (2012).

[20] Ruskey, F. and M. Weston. "A Survey of Venn Diagrams". *The Electronic Journal of Combinatorics*, DS#5 (2005).

[21] Stapleton, G., J. Howse, and P. Rodgers. "A graph theoretic approach to general Euler diagram drawing". *Theoretical Computer Science*, 411 (2010), no. 1 91–112.

[22] Thompson, M. "Venn Polyominoes". URL `http://webhome.cs.uvic.ca/~ruskey/Publications/VennPoly/Thompson/venn_polyominos.html`. Original website no longer available, site copied to a new website by Frank Ruskey.

[23] Venn, J. "On the diagrammatic and mechanical representation of propositions and reasonings". *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 8 (1880) 1–8.

[24] West, D. B. *Introduction to Graph Theory.* Upper Saddle River, NJ: Prentice Hall, second edition (2001).

[25] Winkler, P. "Venn diagrams: some observations and an open problem". *Congressus Numerantium*, 45 (1984) 267–274.