ANALYSIS OF VENN DIAGRAMS USING CYCLES IN GRAPHS

Kiran B. Chilakamarri,
Peter Hamburger,
and Raymond E. Pippert
Indiana University - Purdue University Fort Wayne,
Department of Mathematical Sciences,
2101 Coliseum Boulevard East,
Fort Wayne, IN 46805, U.S.A.
E-mail: hamburge@ipfw.edu
pippert@ipfw.edu

Proposed running head:

ANALYSIS OF VENN DIAGRAMS

Proofs should be sent to:

Peter Hamburger

Indiana University - Purdue University Fort Wayne,
Department of Mathematical Sciences,
2101 Coliseum Boulevard East,
Fort Wayne, IN 46805, U.S.A.
E-mail: hamburge@ipfw.edu

Abstract. This paper is the last in a series by the authors on the use of graph theory to analyze Venn diagrams on few curves (see [1,2,6,7]). We complete the construction (and hence the enumeration) of spherical Venn diagrams on five curves, which yields additional results about conjectures of Grünbaum concerning which Venn diagrams are convex, which are exposed, and which can be drawn with congruent ellipses.

Mathematics Subject Classifications (1991): 05C10, 05C45, 51M04

Key words: Venn diagram, convex Venn diagram, planar graph, Venn graph or dual graph.

1. Introduction.

The motivation for much of our work on Venn diagrams has been problems and conjectures of Grünbaum, who noted in [4], “It is not known how many types of simple Venn diagrams exist with five curves.” This is the basis for the present paper, in which we determine the number of simple Venn diagrams with five curves on the sphere. (All diagrams in the plane can be determined from these, but we choose not to pursue this because the process is laborious and the investigation would yield no further insights into the geometrical and topological properties of the diagrams.) In addition, our investigation allows us to answer other questions posed by Grünbaum about Venn diagrams with special characteristics. Before describing these results and the procedures used to obtain them, we will require some definitions, notations, and preliminary observations.

As the title indicates, we use graph theory for our work; any standard book on graph theory, such as [3], will suffice as a reference for terminology.

An *independent family of curves* in the plane is a collection of simple closed Jordan curves $\mathfrak{F} = \{ C_1, C_2, \dots , C_n \}$ such that each of the 2^n sets $X_1 \cap X_2 \cap \dots \cap X_n$ is nonempty, where X_i is either the bounded interior or the unbounded exterior of C_i for $i = 1,2, \dots ,n$. An independent family of curves is a *Venn diagram* or an *n-Venn diagram* if each of the sets is not only nonempty but also connected. A Venn diagram is *simple* iff at most two curves intersect (transversely) at a point. An *n*-Venn diagram is *reducible* iff there is a curve whose removal leaves an *(n-1)*-Venn diagram. The projection of a Venn diagram from the plane to a sphere by means of stereographic projection yields a *spherical Venn diagram*.

Two Venn diagrams are *isomorphic* if one of them can be changed to the other or its mirror image by continuous transformation of the plane, however, we more often use the concept of graph isomorphism except when considering convex and exposed diagrams, in which case both concepts may arise; we shall be explicit in the use of the term graph isomorphism whenever it is intended. A Venn diagram is said to be *exposed* if each of its curves has an arc on the boundary of the unbounded exterior region; it is *convex* if it is isomorphic to a Venn diagram formed by convex curves.

With each Venn diagram we associate two graphs. The Venn diagram itself can be viewed as a planar graph $V(\mathscr{F})$ where all the intersection points of the curves in $\mathscr{F}$ *are the vertices of* $V(\mathscr{F})$ and the segments of the curves with vertices as end points are the edges of $V(\mathscr{F})$. In proper context, confusion rarely arises from also calling this graph a *Venn diagram*. The graph dual of $V(\mathscr{F})$ will be called the *Venn graph*, denoted by $D(\mathscr{F})$. Several interesting properties of Venn diagrams and Venn graphs were derived in [1]. Here we simply state those properties we need as remarks and refer the reader to [1] for proofs. Throughout this section, we will assume $|\mathscr{F}| = n$ until we begin the analysis of the case $n = 5$.

Remark 1. By the definition of dual graph, it is clear that a curve in $\mathscr{F}$ does not intersect two adjacent edges of a face in $D(\mathscr{F})$.

Remark 2. The number of vertices in an n-Venn graph contained in the intersection of either the interiors or the exteriors of k curves is 2^{n-k}.

Remark 3. Let $V(\mathscr{F})$ be a simple n-Venn diagram. In $V(\mathscr{F})$ there are no 2-faces if $n \geq 3$, and no two edges in any face belong to the same curve. Hence $3 \leq d(x) \leq n$ for every vertex x in $D(\mathscr{F})$, where $d(x)$ is the degree of x.

Remark 4. In a simple Venn graph $D(\mathscr{F})$, every face is a 4-face (quadrilateral), and hence $D(\mathscr{F})$ is a maximal bipartite planar graph.

Let $V(\mathscr{F})$ be a simple n-Venn diagram. If we travel along one of the curves of $\mathscr{F}$, listing the edges of the dual graph $D(\mathscr{F})$ as we encounter them, we note that each consecutive pair of edges in the list is an opposite pair of edges from a face of $D(\mathscr{F})$, and the list of edges forms a matching in $V(\mathscr{F})$ whose size is precisely the number of intersections of the curve. Thus it is clear that the edges of $D(\mathscr{F})$ can be partitioned into matchings $M_1, M_2, \ldots, M_n$.

From the fact that $D(\mathscr{F})$ is a maximal bipartite planar graph, it can easily be shown that each matching has an even number of edges. (This also follows from the fact that the number of transverse intersections of two closed curves must be even.) Since there are 2^n vertices in $D(\mathscr{F})$ and since every curve must intersect every other curve at least twice in the Venn diagram, it follows that $2(n-1) \leq |M_i| \leq 2^{n-1}$ for $i=1,2,\ldots,n$. We denote by x_i the number of matchings having exactly i edges. A curve in $\mathscr{F}$ corresponding to an i-matching will be referred to as an *i-curve.*

Remark 5. In a simple Venn graph $D(\mathscr{F})$, $x_i = 0$ if i is odd or $i > 2^{n-1}$, and

$$\sum x_i = n \qquad \text{and} \qquad \sum i x_i = |E(D(\mathscr{F}))| = 2^{n+1} - 4.$$

Remark 6. Deleting the edges of any of the matchings M_i from a simple Venn graph $D(\mathscr{F})$ disconnects the graph into two subgraphs G_1 and G_2, both having the same number of vertices, edges, and faces. We call G_1 and G_2 *half Venn graphs,* and the corresponding diagrams *half-Venn diagrams.*

We can now summarize the results obtained in this paper.

THEOREM:

i) There are 20 nonisomorphic simple spherical Venn diagrams on five curves (see [6], and Figures 1 and 2).

ii) All simple Venn diagrams on five curves are exposed.

iii) Of the 20 diagrams in i), 11 are convex and 9 are nonconvex (see [6], and Figures 1 and 2).

iv) The 11 convex spherical diagrams in iii) yield 17 nonisomorphic convex Venn diagrams in the plane (see [6], and Figures 2 and 3), of which exactly 7 can be drawn with congruent ellipses (see [6], and Figures 2 and 3).

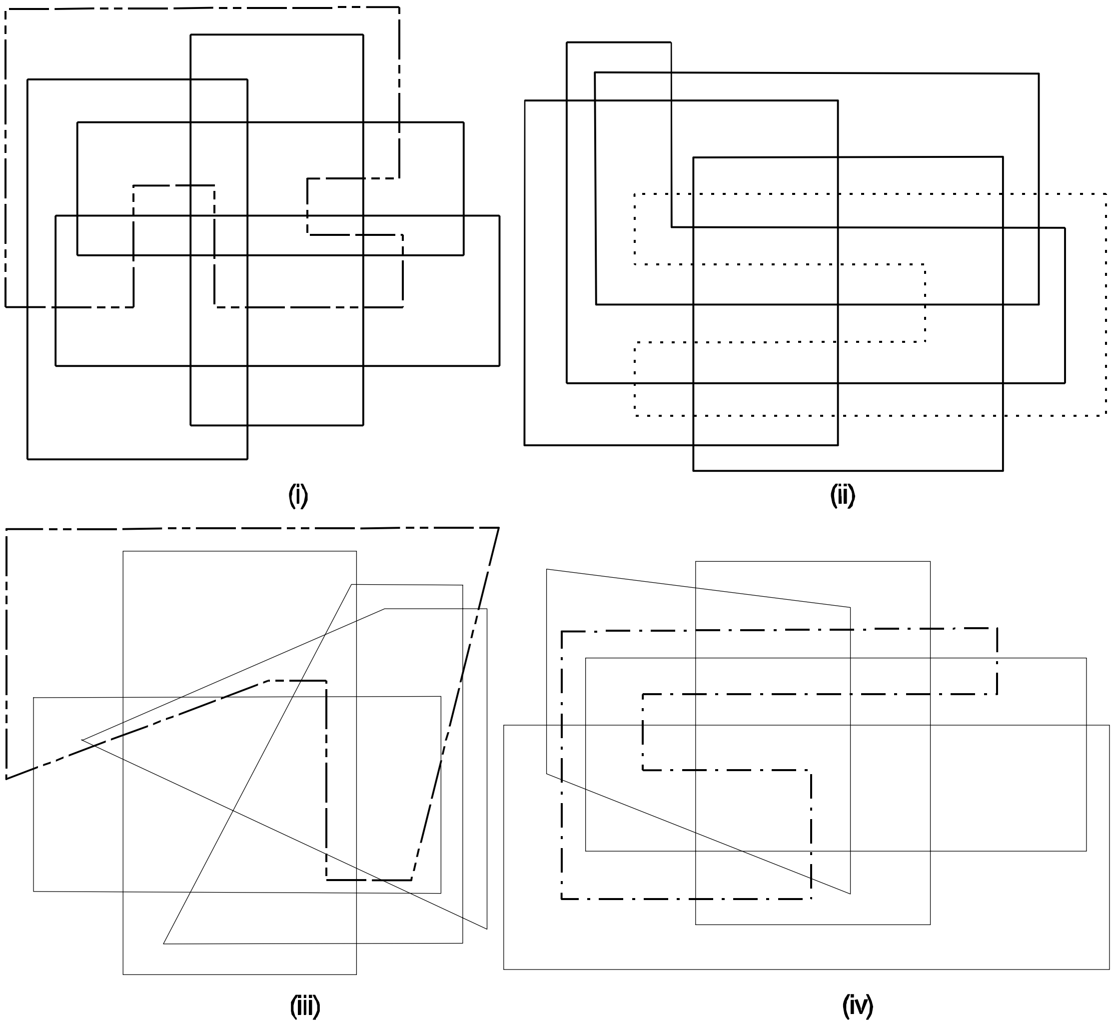
(i) (ii)

(iii) (iv)

Simple, nonconvex, irreducible Venn diagrams

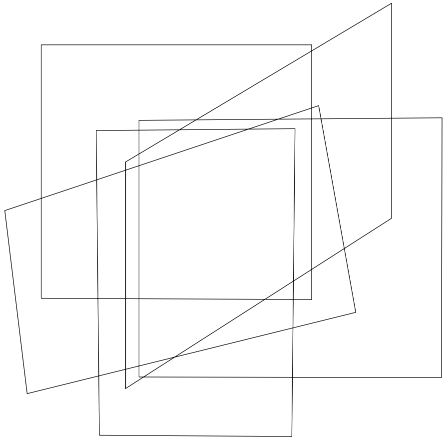

Figure 1

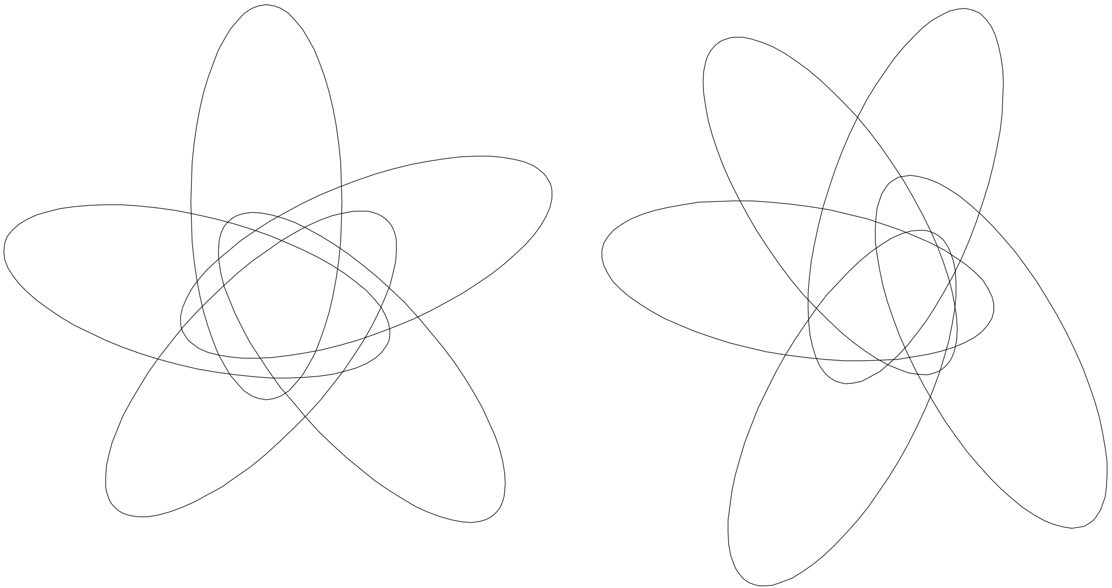

Maximum matching size 12

(i) Maximum matching size 14

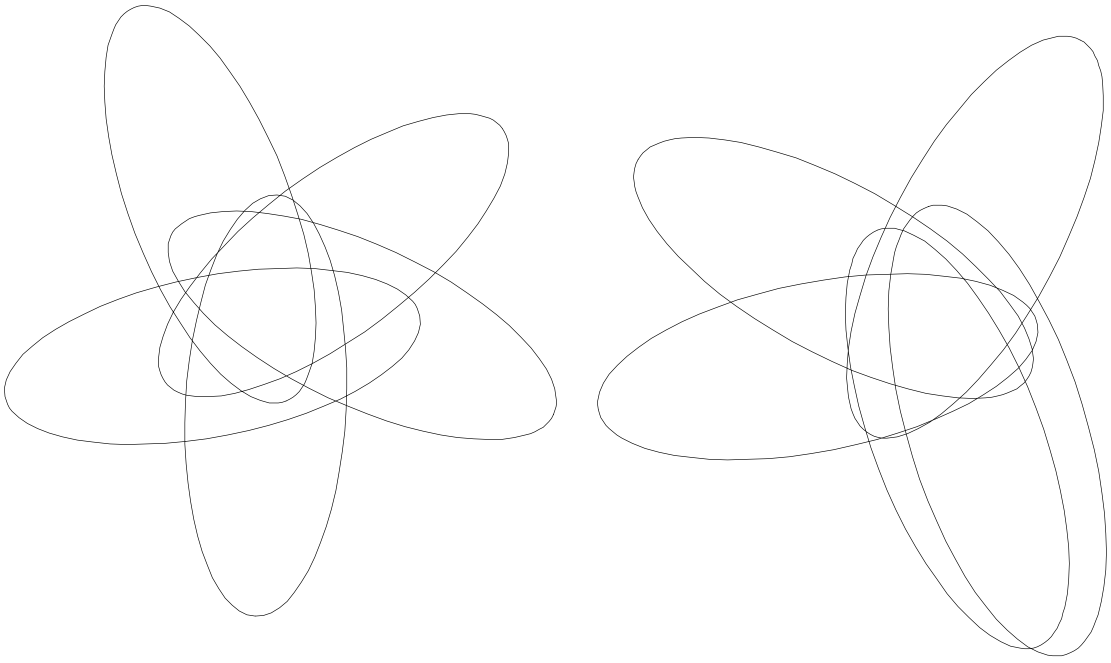

(ii) Maximum matching size 14

Figure 2

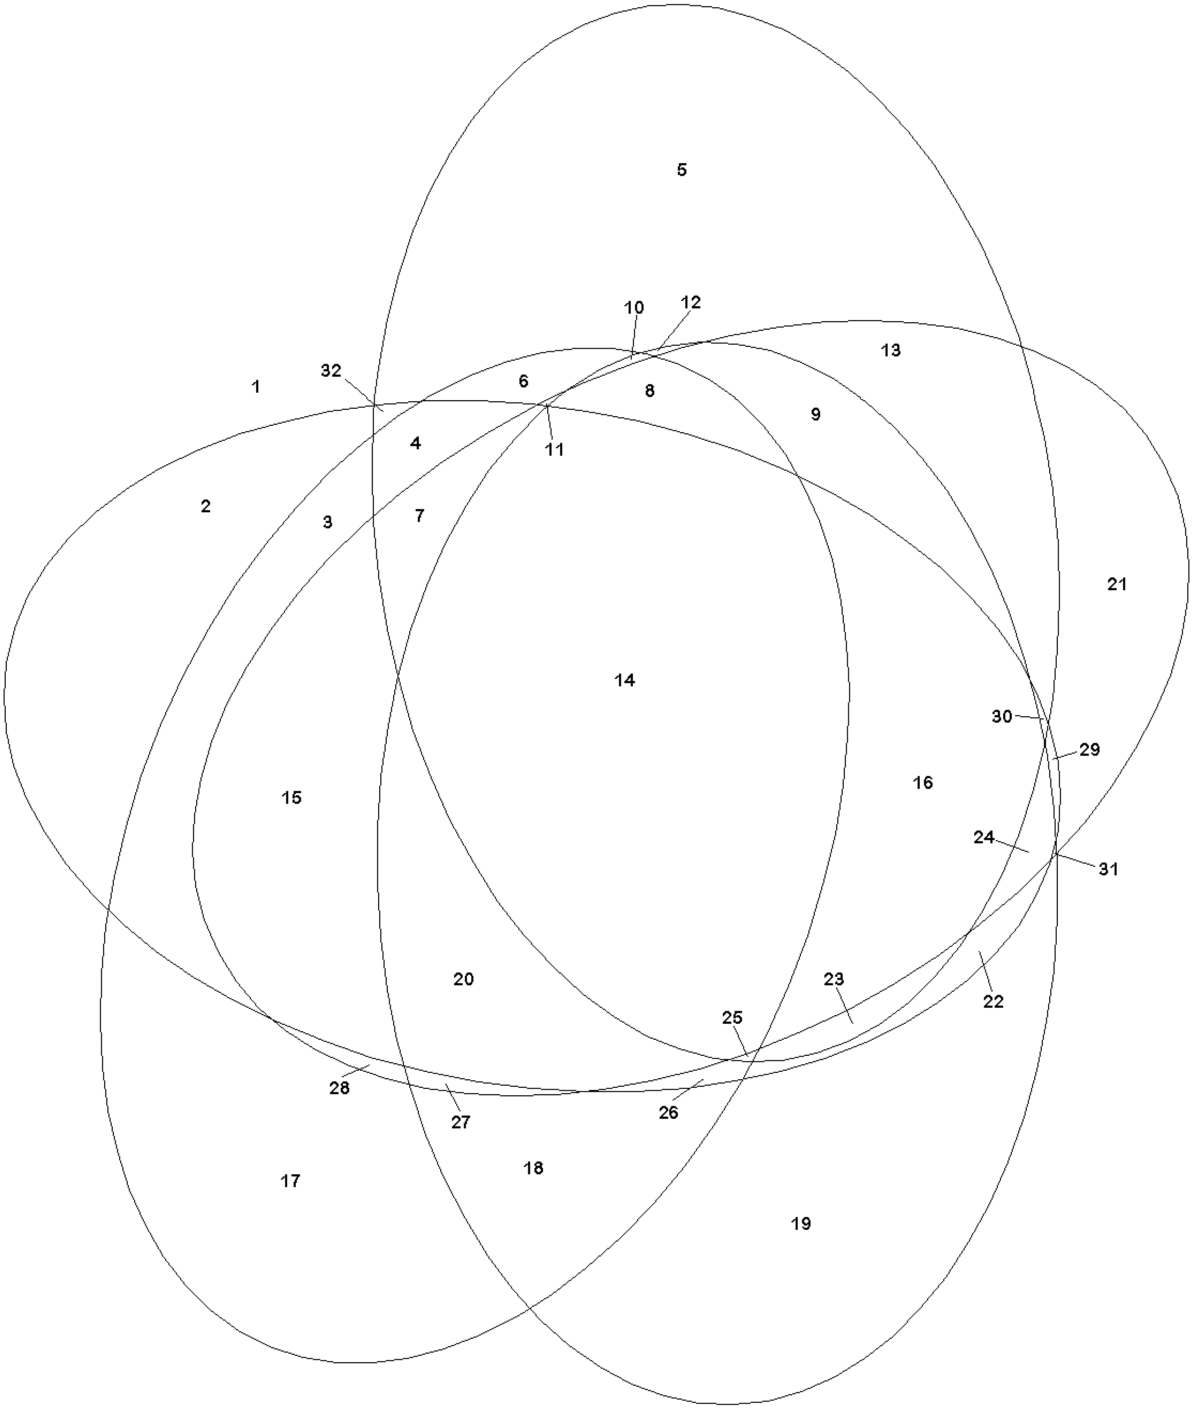
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32

Figure 3

To describe the structure of our investigation, we refer to Remark 4, and note that the distribution of the matching sizes $\{x_i\}$ is not arbitrary. Analysis of possible matching sizes allows a significant reduction in the number of cases consider. Since we are analyzing 5-Venn diagrams, we take n =5, and find that $|V(D(\mathfrak{F}))| = 32$, $|E(D(\mathfrak{F}))| = 60$, and $8 \leq |M_i| \leq 16$ for i =1,2,...,5. From Remark 5, we have

$$x_8 + x_{10} + x_{12} + x_{14} + x_{16} = 5 \qquad \text{and} \qquad 8x_8 + 10x_{10} + 12x_{12} + 14x_{14} + 16x_{16} = 60,$$

where x_i is a nonnegative integer, i = 1,2,...,5. Solving, we obtain the following table of possible matching sizes (with rows numbered for reference).

Row #	x_8	x_{10}	x_{12}	x_{14}	x_{16}
1	2	0	1	0	2
2	1	2	0	0	2
3	2	0	0	2	1
4	1	1	1	1	1
5	0	3	0	1	1
6	1	0	3	0	1
7	0	2	2	0	1
8	1	1	0	3	0
9	1	0	2	2	0
10	0	2	1	2	0
11	0	1	3	1	0
12	0	0	5	0	0

Table 1

A quick study of Table 1 reveals a natural classification of simple Venn diagrams on five curves into three classes depending on whether the size of a maximum matching is 16 (rows 1-7), 14 (rows 8-11), or 12 (row 12).

Venn diagrams having 16-matchings are constructed by finding all distinct Hamilton cycles in the dual graph [6]; we include a brief discussion of this case in Section 2 for purposes of completeness. In Section 3, we obtain diagrams having maximum matching size 14 by determining a maximal substructure common to all such diagrams and then extending them to find cycles having the needed properties. A similar approach was used to simplify the proof of a theorem (5.1) in [8] (see [9]). In Section 4, we consider the case in which the maximum matching has size 12; the procedure we use is to transform the problem of enumerating Venn diagrams into one of enumerating maximal bipartite planar graphs of a certain type. This method was first used in [1] to prove that there is a unique simple spherical Venn diagram on four curves (a result mentioned earlier by Grünbaum).

2. Maximum matching size 16

This case was analyzed completely in [6]. The Venn diagrams in this case are reducible diagrams, since the deletion from a 5-Venn diagram of any curve corresponding to a 16-matching yields a 4-Venn diagram. Thus all 5-Venn diagrams in this case can be constructed as extensions of some 4-Venn diagram. But there is only one simple spherical 4-Venn diagram [1] which we denote by *V*(4), with Venn graph *D*(4). The extensions can be listed by finding all Hamilton cycles in *D*(4) and adding them to *V*(4) as a fifth curve, thus obtaining 5-Venn diagrams. The results are that there are 11 distinct simple, reducible spherical 5-Venn diagrams, each of which has a representation in the plane as an exposed diagram, but five of which cannot be represented as a convex one.

3. Maximum matching size 14.

We begin this case by observing that the deletion of any 14-curve must result in an independent family of curves having 18 regions, since the curve intersects 14, making the total 32 as required. Two pairs of the 18 regions must each correspond to a single set determined by C_1, C_2, C_3, and C_4. To obtain a simple Venn diagram, the 14-curve must separate each of these pairs and intersect every other region.

By a sequence of lemmas, we show that in any simple 5-Venn diagram having maximum matching size 14, there is a 14-curve and another curve whose deletion yields a simple 3-Venn diagram (which is unique -- see [1]). This will provide the basis for our construction, since it allows us to begin with a known structure.

Lemma 1. In an irreducible simple 5-Venn diagram, no pair of curves C_i and C_j have eight intersections.

Proof. Remarks 1 and 3 show that the maximum matching would have to be larger than 14 for this to occur. ▌

Lemma 2. Three curves from a simple 5-Venn diagram form a simple sub-3-Venn diagram if and only if they have a total of six intersections.

Proof. Since the simple 3-Venn diagram is unique, a simple sub-3-Venn diagram clearly has six intersections. Now suppose three curves from a simple 5-Venn diagram have six intersections. Since each curve intersects each other curve at least twice, these curves must intersect each other exactly twice. That means that any pair is a 2-Venn diagram. Since the three curves must form an independent family, the third curve must intersect all four regions of the 2-Venn diagram, but can intersect each curve only twice. This forces it to be the unique simple 3-Venn diagram. ▌

Lemma 3. A simple 5-Venn diagram with maximum matching 12 contains no simple sub-3-Venn diagram.

Proof. The removal of two curves leaves the fewest remaining intersections if the two curves intersect each other minimally, i.e., twice, thus allowing the maximum number of intersections with the remaining curves. This forces the number of intersections among the three remaining curves to be at least eight (30 - (2+10+10)). By Lemma 2, these curves cannot form a simple 3-Venn diagram. ∎

Lemma 4. A simple 5-Venn diagram with maximum matching 14 contains a simple sub-3-Venn diagram.

Proof. We proceed by a series of sub lemmas, followed by an analysis of the pertinent cases from Table 1.

Sub Lemma 4.1. If there are two 14-curves, they cannot intersect only twice.

Proof. The removal of two such curves leaves three curves with a total of only four intersections, which is impossible.

Sub Lemma 4.2. If there are either two 14-curves that intersect four times or a 14-curve and a 12-curve that intersect twice, then their removal results in a simple 3-Venn diagram.

Proof. In each case, the removal of the two curves leaves three curves with six intersections, and the result follows from Lemma 2.

Sub Lemma 4.3. If there are three 14-curves, then two of these intersect four times, so their removal leaves a simple 3-Venn diagram.

Proof. Suppose there is no pair with four intersections. Since Sub Lemma 4.1 precludes any pair with two, each pair must have at least six intersections. This leaves too few for the 14-curves to intersect with the remaining curves. Hence there must be a pair with four intersections.

Now from Table 1, we see that there are four cases to consider in which the maximum matching size is 14, occurring in rows 8-11 of Table 1.

Row 8 is the only case with three 14-curves. It is finished by Sub Lemma 4.3.

In row 9, either the two 14-curves have four intersections and we are done by Sub Lemma 4.2, or else Lemma 4.1 implies that they intersect six times. Then one of the 14-curves and a 12-curve must intersect twice, since otherwise the 14-curves would already have 14 intersections with each other and the 12-curves. In either case, there is a simple sub-3-Venn diagram.

Row 10 is finished by Sub Lemma 4.2 if either the two 14-curves intersect four times, or they intersect six times and one of them intersects the 12-curve twice. This leaves only the alternative that the intersection table is (except for a possible relabeling of the curves) that of Table 2, where

C_1 and C_2 are 14-curves and C_3 is the 12-curve.

	C_1	C_2	C_3	C_4	C_5
C_1	0	6	4	2	2
C_2	6	0	4	2	2
C_3	4	4	0	2	2
C_4	4	2	2	0	2
C_5	2	2	2	4	0

Table 2

Now C_4 and C_5 intersect four times, while C_3 intersects each of them twice. We can draw C_4 and C_5 as in Figure 4.

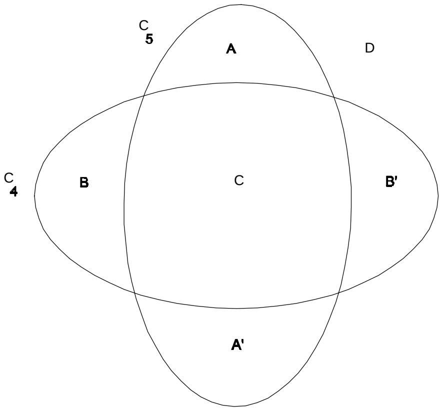

Figure 4

Now C_3 must divide both regions C and D. The constraints of simplicity plus only two intersections with each of the other curves require that (up to symmetry) it pass through A and A' or through A and B'. In the first case, we obtain the solid curves in Figure 5, and C_2 must follow the dashed curve.

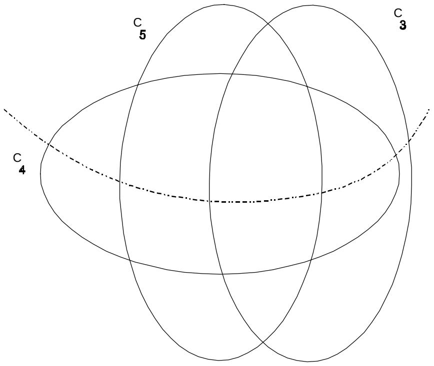

Figure 5

Since C_2 cannot re-enter either C_4 or C_5, we are unable to obtain the 18 regions needed, as observed at the beginning of this section.

In the second case, we obtain the solid curves in Figure 6. Now a portion of A must still be separated from A', and similarly for B and B'. However, to separate A from A', a curve must pass through B, and similarly a curve must pass through A'. These must be different curves, because a curve passing through a region cannot separate it from another region. (One part is in the interior and another in the exterior.) Thus, C_2 must divide one of the 2-faces A' or B as shown by the dashed curve. But it must also divide each of E, F, G, and H, which is impossible with the constraint that it intersect each of C_4 and C_5 only twice.

Thus all possible simple 5-Venn diagrams from Row 10 contain a simple sub-3-Venn diagram.

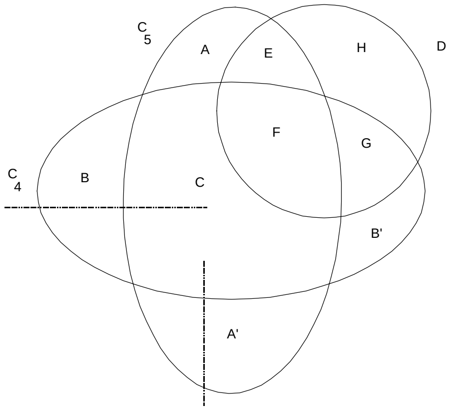

Figure 6

Row 11 is finished if the 14-curve intersects any 12-curve twice; the only alternative is that the intersection table is (again, up to relabeling) that of Table 3.

	C_1	C_2	C_3	C_4	C_5
C_1	0	4	4	4	2
C_2	4	0	2	2	4
C_3	4	2	0	4	2
C_4	4	2	4	0	2
C_5	2	4	2	2	0

Table 3

We note that the argument in the case of Row 10 depended only on the curves C_2 through C_5, which yield the same sub-table as that in Table 3, with appropriate relabeling. The same argument then shows that this table cannot yield a simple 5-Venn diagram.

This completes the proof of Lemma 4. ∎

Now we proceed with the construction of all nonisomorphic independent families of four curves on the sphere containing a 3-Venn diagram and 18 regions as described above. For brevity, we refer to these as *4-sets*. Although we are working on the sphere, we are naturally drawing diagrams in the plane; without loss of generality, we may thus assume that one of the pairs of regions which must be separated are adjacent with the outer face. Since there are no 2-faces in the diagram, by symmetry, there are only two ways this can occur, as shown in Figure 7.

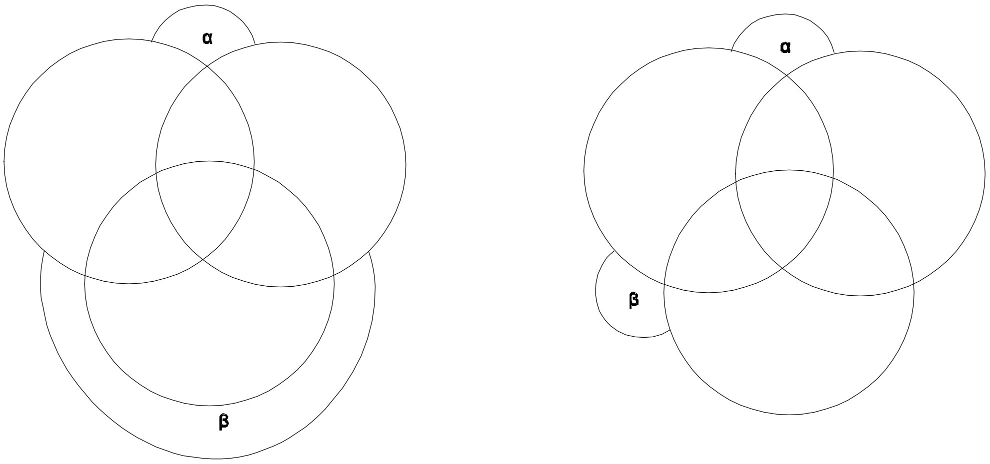

(a) (b)

Figure 7

Our constructions are done in two cases, according as we begin with Figure 7 (a) or (b).

Case 1. In Figure 7 (a), if the extension of one end of the α-curve enters any of regions C, E, F, or G, the other cannot, because it could not leave without creating too many new regions. Thus, up to symmetry, the only 4-set is that in Figure 8.

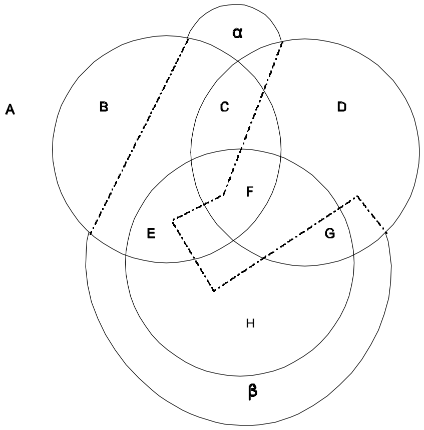

Figure 8

Case 2. There are numerous possibilities in this case. We note that region D is symmetric with H, and C is symmetric with E.

We first claim that if the α-curve and the β-curve are connected in B (without leaving B), then the other pair (γ,δ) of regions to be separated by the fifth curve cannot be in B, C, E, or G. One of a pair in B would be a 2-face which must then be both divided and separated by C_5, which we have observed is impossible. Since the curve can't re-enter B, creating a pair in C (or E) would require the creation of another pair in an adjacent region, resulting in too many regions. Similarly, creating a pair in G also results in too many regions. Thus, the other pair must lie either in D (or H) or in F.

Subcase 2.A. C_4 is connected in B.

i) The pair (γ,δ) lies in D. There are two ways this can occur.

a) Regions γ and δ are as in Figure 9, with the resulting 4-set as shown.

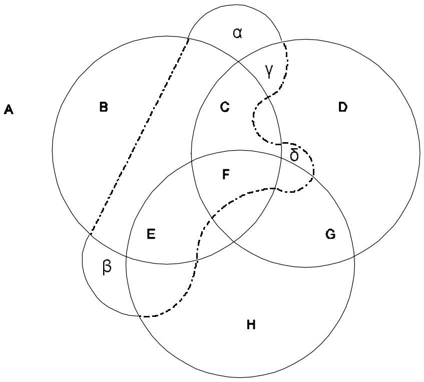

Figure 9

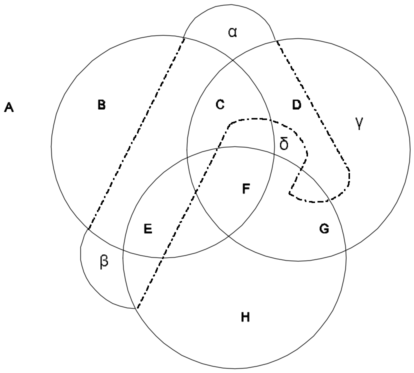

Figure 10

b) Regions γ and δ are as in Figure 10, with the resulting 4-set as shown.

ii) The pair (γ,δ) lies in F. C_4 must cut corners (i.e., create 3-faces), for otherwise too many regions are created. Up to symmetry, there are again two cases.

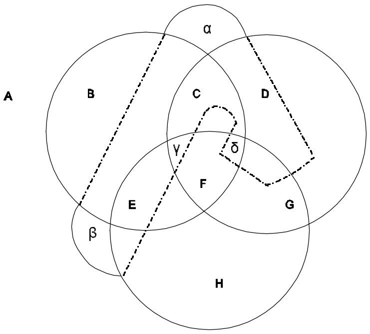

Figure 11

a) Regions γ and δ are as in Figure 11, with the resulting 4-set as shown.

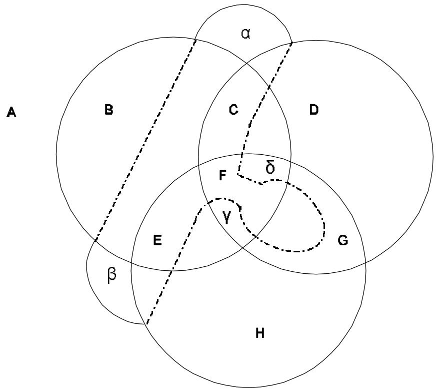

Figure 12

b) Regions γ and δ are as in Figure 12, with the resulting 4-set as shown.

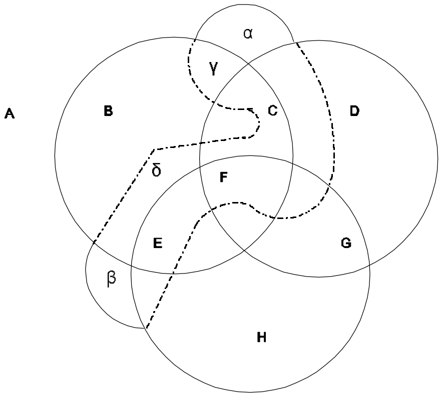

Figure 13

Subcase 2.B. C_4 is not connected in B (C_4 leaves B and returns).

i) C_4 enters and leaves C (or E). The result is that of Figure 13, and is isomorphic to Case 1 (make B the exterior region).

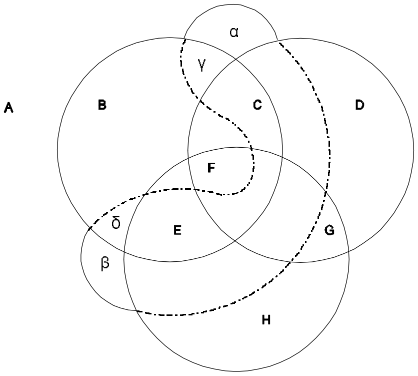

Figure 14

ii) C_4 enters one region and leaves another, as shown in Figure 14.

The seven figures of Cases 1 and 2 give us six different 4-sets, up to symmetry, from which to attempt to construct simple 5-Venn diagrams.

We first construct the dual graph for each of the independent sets of four curves, then find all 14-cycles which separate the regions α from β and γ from δ, as required. These 14-cycles constitute the fifth curve in a simple 5-Venn diagram. The 14-cycles cannot pass through any of α, β, γ, or δ, but in each case, there are additional constraints that help to reduce the number of possibilities. Edges that are forced to be in the cycle are shown bold, while those that must be excluded are dashed. The cases follow those above.

Case 1. The existence of a vertex of degree two and the separation requirements result in the dual graph of Figure 15.

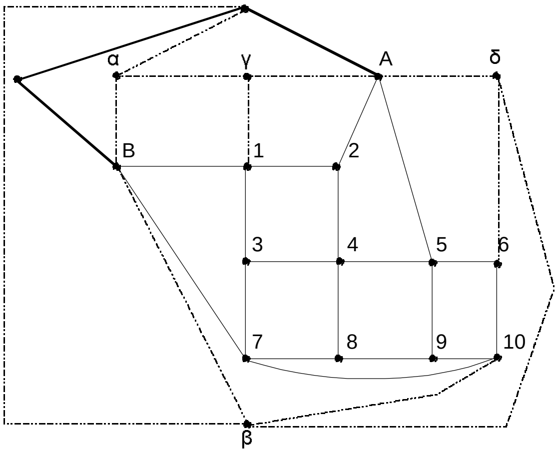

Figure 15

We begin at A and end at B. There are two choices at A: 2 or 5. From 2 we can go to 1 or 4, and from 4, we have the choice of 5 or 8. If we choose 5 first, we have no choices until we reach 8. The possible 14-cycles can be easily traced on the diagram, which we number for reference:

1. A,2,1,3,4,5,6,10,9,8,7,B
2. A,2,1,3,4,8,9,5,6,10,7,B
3. A,2,4,5,6,10,9,8,7,3,1,B
4. A,2,4,8,9,5,6,10,7,3,1,B
5. A,5,6,10,9,8,4,2,1,3,7,B
6. A,5,6,10,9,8,7,3,4,2,1,B

Note that C_4 intersects edges {1,2}, {3,4}, {5,6}, {7,8}, {9,10}, and the path {A,B}, so the constraint that the maximum matching number is 14 allows the use of only four of these edges in C_5. The cycles 1 and 6 each use six of the edges listed above, so they are considered in the 16-case. The Venn diagrams obtained from cycles 2, 3, 4, and 5 are seen in Figures 2 (iii), 1 (i) and (ii), and 2 (ii), respectively.

Case 2.A.i.a. The dual graph, shown in Figure 16, yields only two possibilities for starting at A and ending at B.

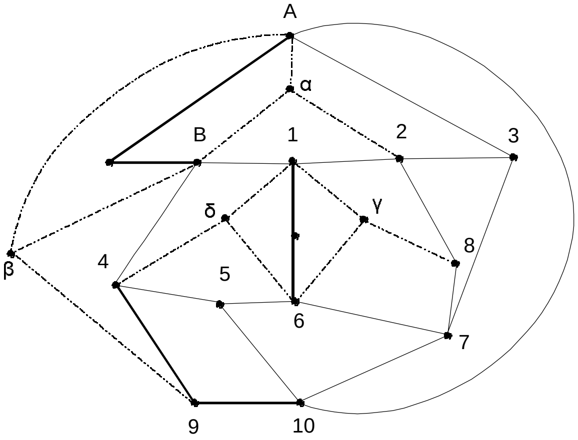

Figure 16

The paths yielding the fifth curve are:

1. A,5,4,3,8,9,13,12,11,10,6,7,B
2. A,10,6,7,11,12,8,9,13,5,4,3,B

Note that C_4 intersects edges {4,5}, {6,7}, {8,9}, {10,11}, {12,13}, and {B,1}, so as in Case 1, the constraint that the maximum matching number is 14 allows the use of only four of these edges in C_5. Both cycles 1 and 2 use six of the edges listed above, so they are considered in the 16-case.

Case 2.A.i.b. In this case, many paths are forced because of degree two vertices, as shown in Figure 17.

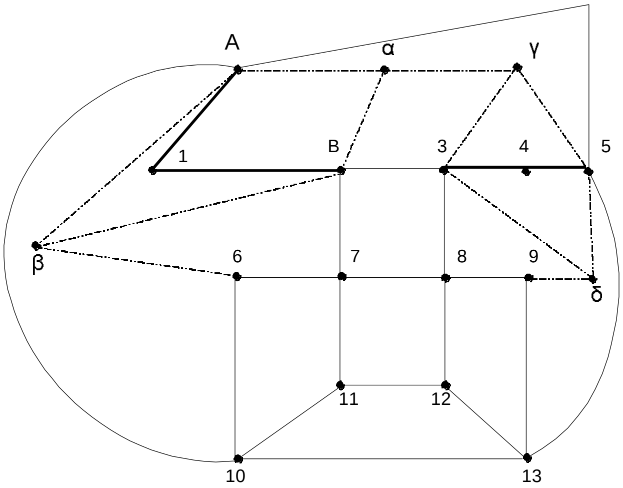

Figure 17

The fact that 1,2,3 is forced means that vertex 4 cannot be included in the desired 14-cycle, so this case yields no potential fifth curve.

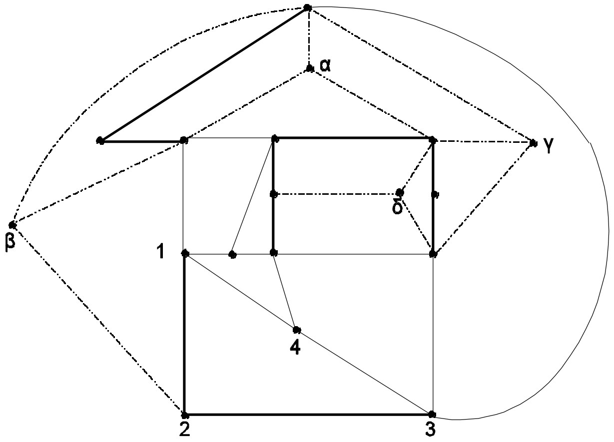

Figure 18

Case 2.A.ii.a. The positions of the vertices α, β, γ, and δ, and the corresponding edges which cannot be in the fifth curve force many edges to be in the curve, as shown in Figure 18.

This allows only two possible 14-cycles of the desired type:

1. A,3,2,8,7,10,9,4,5,6,1,B
2. A,3,7,8,2,1,6,5,10,9,4,B

Note that C_4 intersects edges {2,3}, {4,5} {7,8}, {1,6}, {9,10}, and the path {A,B} so the constraint that the maximum matching number is 14 allows the use of only four of the edges listed above in C_5. The cycle 1 uses six of them, so this is considered in the 16-case. The Venn diagram obtained from cycle 2 is seen in Figure 2 (i).

Case 2.A.ii.b. In the dual graph of Figure 12 it is easily seen that the constraints force a 4-cycle, so no Venn diagrams result from this case.

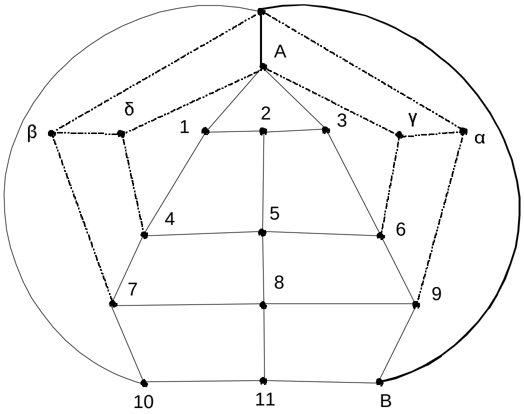

Figure 19

Case 2.B.i. As observed in the construction of the 4-sets, this is isomorphic with Case 1, so no new Venn diagrams can be obtained from this case.

Case 2.B.ii. This is a highly symmetric case, with many 14-cycles to be considered. In Figure 19 we observe that there is both a lateral symmetry, a front-to-back symmetry (viewing the diagram as 3-dimensional), and the combination of the two. We also note that C_4 intersects edges {1,4}, {2,5}, {3,6}, {7,10}, {8,11}, and {9,B}, so once again the constraint that the maximum matching number is 14 allows the use of only four of these edges in C_5. Because of the lateral symmetry, we can choose one of the side edges to be in C_5; we select the heavy edge incident with B, leaving only the front-to-back and combination symmetries. The results are that we obtain eight 14-cycles:

1. A,1,2,3,6,5,4,7,10,11,8,9,B
2. A,1,2,3,6,9,8,5,4,7,10,11,B
3. A,1,2,5,4,3,6,9,8,7,10,11,B

4. A,1,4,7,10,11,8,5,2,3,6,9,B
5. A,3,2,1,4,5,6,9,8,7,10,11,B
6. A,3,2,1,4,7,10,11,8,5,6,9,B
7. A,3,6,5,2,1,4,7,10,11,8,9,B
8. A,3,6,9,8,5,2,1,4,7,10,11,B

Now, front-to-back symmetry shows that cycles 6 and 8 are the same, while the combination symmetry shows 1 and 3 to be the same. Furthermore, cycles 4 and 7 each use six of the edges listed above, so they are considered in the 16-case. Consequently, up to symmetry there are only four 14-cycles which yield Venn diagrams in this case:

1. A,1,2,3,6,5,4,7,10,11,8,9,B
2. A,1,2,3,6,9,8,5,4,7,10,11,B
5. A,3,2,1,4,5,6,9,8,7,10,11,B
6. A,3,2,1,4,7,10,11,8,5,6,9,B

The Venn diagram obtained from cycle 1 is isomorphic to the Venn diagram that is obtained from cycle 3 of Case 1.

The Venn diagrams obtained from cycles 2, 5, and 6 can be seen in Figure 1 (ii) and (iii), and the convex unlabeled one, respectively.

Finally, that the diagrams in Figures 1 and 2 are not isomorphic can be checked by counting the intersection numbers of all the curves, the number of 5-faces, and considering the relative position of the 5-faces. This is tedious but straightforward and is left for the reader.

We summarize the results of Section 3.

PROPOSITION:

There are eight spherical Venn diagrams having maximum intersection number 14.

4. Maximum matching size 12.

We proceed either by constructing half-Venn graphs which will be joined in pairs to form Venn graphs which then yield the Venn Diagrams, or by constructing half-Venn diagrams to be joined to form Venn diagrams directly. As mentioned earlier, we consider spherical diagrams unless otherwise noted. When the maximum matching size is twelve, we use the intersection numbers μ_{ij} ($1 \leq i,j \leq 5$) to restrict the number of cases that must be considered. From Table 1, we see that all of the matchings must have size twelve, and since every pair of curves intersects at least twice, we have $4 \leq \max(\mu_{ij}) \leq 6$ (of course μ_{ij} is even). The matrix $[\mu_{ij}]$ is symmetric with zeroes on the diagonal, and each row and column adds to twelve.

Case 1. Max $\mu_{ij} = 6$.

Without loss of generality we take μ_{12} to be 6 and assume that the half-Venn graphs are joined along the matching corresponding to C_1. Then the curve C_2 can intersect C_1 in two possible ways as shown in Figure 20. We shall use portions of the potential Venn diagram to eliminate some candidates for half-Venn diagrams, giving more detail in this first case than will be provided in the rest.

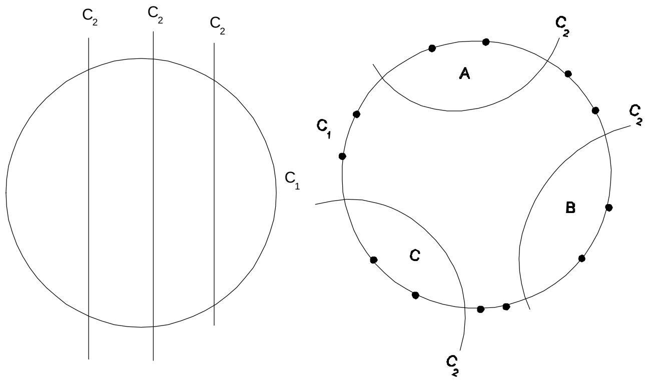

(a) (b)

Figure 20

Subcase 1.1 C_1 and C_2 intersect as in Figure 20 (a)

In Figure 20 (a), there is only one way to connect the segments of C_2, as shown in Figure 21.

Observe that there are four 2-faces, labeled A, B, C, and D, none of which are allowed in a Venn diagram as noted in Remark 3; consequently, another curve must enter each of these regions. If it crosses C_1 to enter one region, it cannot cross C_1 to leave the region because each curve has exactly twelve intersections and C_2 already intersects C_1 six times, leaving only two for each of the others; thus if the new curve enters and leaves by crossing C_1, it cannot cross again, so it cannot reach the other regions that must be divided to maintain an independent family. The same argument can be applied to C_2. Consequently, the curve must enter the region by crossing one curve, and leave by crossing the other.

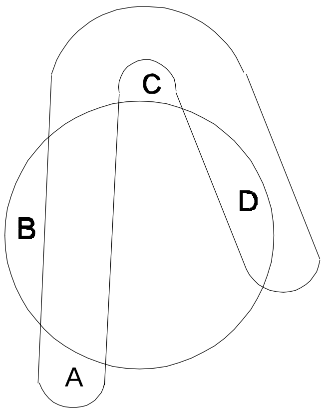

Figure 21

Let C_3 be the curve passing through region C in Figure 21. When it leaves C_1, it must intersect one of the three remaining segments, and the constraints imposed by the intersection numbers are readily seen to allow only the three possibilities shown in Figure 22.

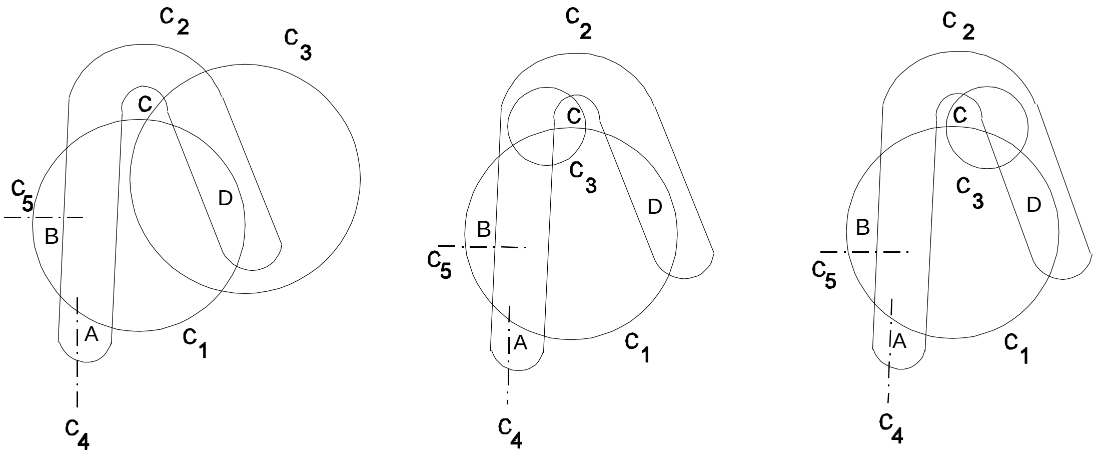

(a) (b) (c)

Figure 22

We now consider the remaining 2-faces. In each case, the segments intersecting A and B cannot be parts of the same curve, because they would have already crossed both C_1 and C_2 twice, so they would have to be connected both in $C_1 \cap C_2$ and outside $C_1 \cap C_2$, yielding as above a curve that cannot intersect the other regions that must be divided to maintain an independent family. We label the curve through A as C_4, and the one through B as C_5. Now in each of Figures 22 (a) and (b), the curve through D must be either C_4 or C_5. In either case, when it has passed through D, that curve already has two intersections with each of C_1 and C_2, so it cannot be closed, *i.e.*, these cases are not candidates for Venn diagrams. In Figure 22 (c), C_5 cannot reach all of the regions inside C_3, cross each of C_1 and C_2 exactly once, and close itself, so it cannot yield a Venn diagram either. We conclude that the configuration of Figure 22 (c) yields no half-Venn diagrams.

Subcase 1.2 C_1 and C_2 intersect as in Figure 20 (b)

We denote the interior and exterior of a closed Jordan curve C by ${}^{I}C$ and ${}^{O}C$, respectively. We now concentrate on the construction of half-Venn graphs, and consider the interior of C_1. There are twelve vertices of $D(\mathfrak{I})$ in ${}^{I}C_1$ from the matching M_1 corresponding to C_1, so there must be four other vertices in ${}^{I}C_1$, since there are sixteen vertices of the 5-Venn graph inside C_1 and sixteen outside. By similar reasoning, since the regions comprising ${}^{I}C_1 \cap {}^{I}C_2$ already contain six vertices from M_1, exactly two of the four must be inside the regions and two outside.

Now suppose that there is exactly one new vertex *x* in one of the regions of ${}^{I}C_1 \cap {}^{I}C_2$, say A. Since a Venn graph contains no odd cycles, *x* can be adjacent to at most one vertex of M_1, and from Remark

1 there can be at most one edge incident with *x* that crosses C_2. This implies that d(*x*) = 2 in contradiction to the information in Remark 3. We conclude that both of the new vertices, *x* and *y*, in ${}^{I}C_1 \cap {}^{I}C_2$ lie in a single region A. Degree constraints show immediately that *x*, *y*, and the two vertices from M_1 in A must form a 4-cycle, while each of *x* and *y* must be incident with an edge which crosses C_2.

Since there are no new vertices in any of the remaining regions of ${}^{I}C_1 \cap {}^{I}C_2$, {d,g} and {h,r} must be edges of $D(\mathfrak{S})$ (see Remark 1). The vertex *y* is adjacent with exactly one more vertex from the set {r,g,u,v} since any other choice produces an odd cycle or forces d(*x*) = 2. We consider the possibilities in order. Suppose *y* is adjacent to r. Then d(r)=5 (we say r is *saturated*, since from Remark 3 it can have no more edges). Since every face of $D(\mathfrak{S})$ is a 4-cycle, the adjacency of h and x is forced, saturating h, and making it impossible to complete $D(\mathfrak{S})$. Thus y is not adjacent to r, so suppose it is adjacent to g. Then g is saturated and no new edges can be incident with any of g, y, or a. This means they cannot be incorporated into a 4-cycle, so again we cannot obtain $D(\mathfrak{S})$. Thus y is not adjacent to g. The symmetry of x and y allows us to suppose that x is adjacent to u and y is adjacent to v. Remarks 1 and 2 imply that {s,v}, {c,u}, and {u,v} are all edges of $D(\mathfrak{S})$ as shown in Figure 23.

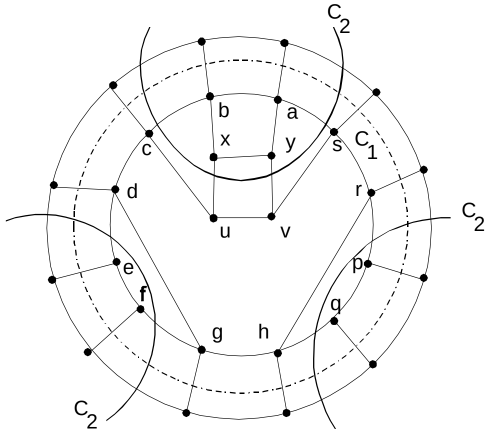

Figure 23

In order to convert the 8-cycle u,v,s,r,h,g,d,c,u into 4-cycles, we must add two more edges. We note that each vertex except u and v becomes saturated with the addition of one more incident edge. The symmetry of the existing structure allows us to examine only one of u and v, say u. The fact that we must obtain 4-cycles restricts additional neighbors of u to the set {g,r}. If u is adjacent to both g and r, we have a possible half-Venn graph (Figure 24 (a)). If u is adjacent to g but not r, the saturation of g allows only the additional edge {v,h}, which also yields a possible half-Venn graph (Figure 24 (b)). If u is adjacent to r but not g, then the edge {c,h} is forced (Figure 24 (c)). The case where neither u nor v has additional edges is precluded by the fact that adding the edge {c,s} produces a matching corresponding to a curve in the potential Venn diagram with only four intersections (all curves must have twelve). Thus, the only possible half-Venn graphs from Case 1 are the three in Figure 24.

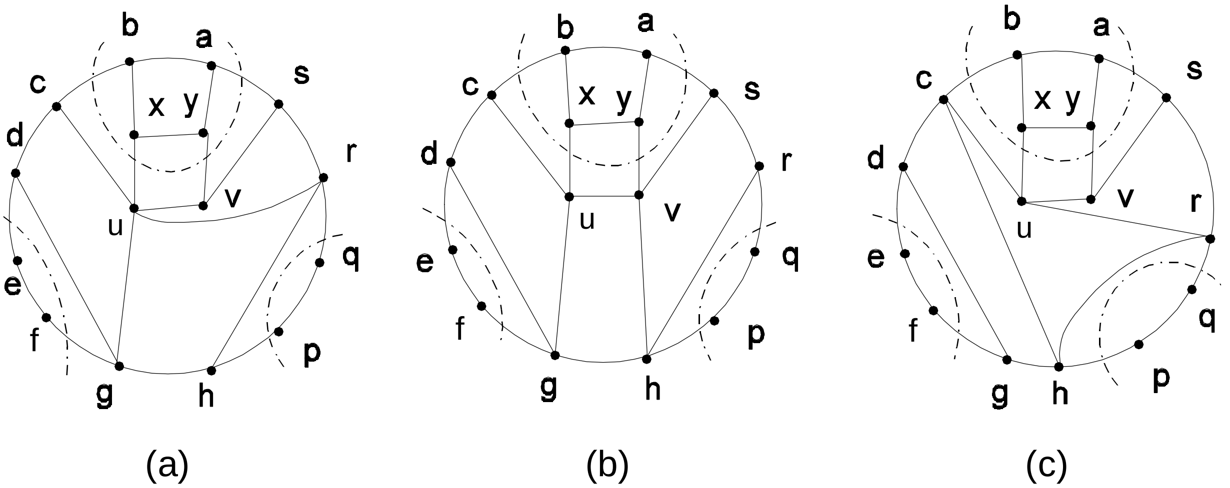

Figure 24

The graph of Figure 24 (a) cannot be a half-Venn graph because a curve, say C_3, crossing C_1 between g and h encloses only six rather than eight vertices of the proposed half-Venn graph when it has exited, and it cannot re-enter. The graph of Figure 24 (c) also cannot be a half-Venn graph because the curve crossing between e and f is self-intersecting.

To analyze the graph of Figure 24 (b), call it G and take another copy of it, say G' with corresponding labels a', b', ... To create a Venn graph from these half-Venn graphs, G and G' must be combined so that the segments of C_2 connect to form a closed Jordan curve. This allows only two alternatives (up to symmetry): match the vertices c, b, a, s with a', s', r', q' in that order, or with p', h', g', f'. The first matching results in a self-intersecting curve, so we discard it. The second alternative does indeed yield a Venn diagram, however, it is easily seen that the curve passing through the graph vertically has only eight intersections, which is not relevant to the case under consideration.

From our results so far, we see that Case 1 yields no Venn diagrams of the desired type.

Case 2. Max $\mu_{ij} = 4$.

We begin this case by constructing the matrix $[\mu_{ij}]$. Without loss of generality, we may choose the entries for C_1 as shown in Table 4; this determines the rest of the table, except for the possible relabeling of C_4 and C_5.

	C_1	C_2	C_3	C_4	C_5
C_1	0	4	4	2	2
C_2	4	0	2	4	2
C_3	4	2	0	2	4
C_4	2	4	2	0	4
C_5	2	2	4	4	0

Table 4

We note from the table that for any curve C_i there are exactly two curves C_j and C_k such that $\mu_{ij} = \mu_{ik} = 2$ and $\mu_{jk} = 4$. This provides us some information about structure. Without loss of generality, we let $i = 1$, $j = 4$, and $k = 5$. As before, we work from the matching M_1 corresponding to the curve C_1. Since $\mu_{14} = \mu_{15} = 2$, the curves C_4 and C_5 each enter and exit the interior of C_1 exactly once; but $\mu_{45} = 4$, which allows three types of structures, depending on the number of intersections of C_4 and C_5 in the interior of C_1 as shown in Figure 25.

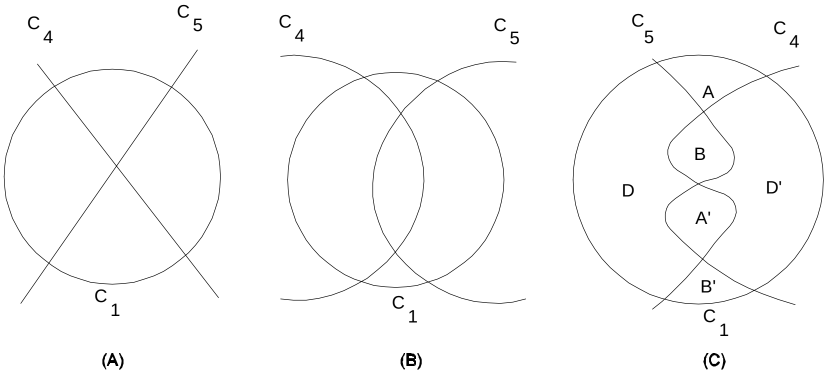

Figure 25

Clearly the intersection numbers force half-Venn diagrams of type A in Figure 25 to be paired with those of type C, while those of type B must be paired with another of the same type. Accordingly, we consider two subcases.

Subcase 2.1 Half-Venn diagrams of types A and C.

In the half-Venn diagram of type C, regions A and A’ have the same description and hence must be separated by another curve; similarly B and B’ must be separated. By symmetry, we assume C_2 separates A from A’, i.e., one set is in the interior of C_2 and the other is in the exterior; and C_3 separates B from B’.

Now μ_{12} and μ_{13} are both four, so there must be eight more intersections on C_1. Since C_2 separates A from A’, it cannot enter region A, and consequently C_3 can enter A at most once, yielding one possible intersection. Similarly, C_2 can enter B’ at most once. But there can be at most three intersections in either D or D’ for otherwise there would be too many regions there. This forces the intersections to be distributed as shown in Figure 26 (a). Any crossing in D or D' creates too many regions, so there are just two possible configurations, only one of which is shown in Figure 26 (b) because of the symmetry.

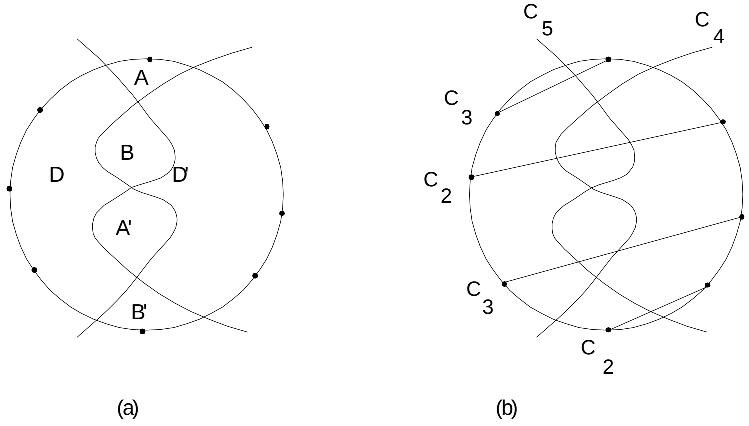

Figure 26

The corresponding half-Venn diagram of type A is shown in Figure 27 (a).

By symmetry, we may choose C_2 to be the curve passing through vertex a. Now vertex a cannot be connected to vertex b (or h), because matching this half-Venn diagram with one of type C would either connect two portions of different curves or make a short cycle; it cannot be connected to vertex d (or f) because this would either connect two portions of different curves or would require that it be self-intersecting; and it cannot be connected to vertex e because again this would connect portions of two different curves in the type C half-Venn diagram. We conclude that vertex a must be connected to vertex c (or g by symmetry). Similarly, we determine that vertex e must be connected to vertex g by C_3. To avoid connecting portions of different curves in the type C half-Venn diagram, we obtain the unique (up to symmetry) configuration in Figure 27 (b). This can be paired with a type C half-Venn diagram to obtain a 5-Venn diagram, but it does not conform to Table 4 and does not have maximum matching 12. Consequently, subcase 2.1 yields no Venn diagrams.

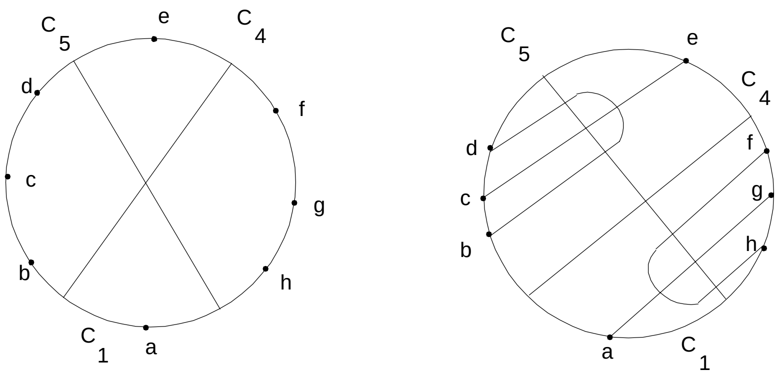

(a) (b)

Figure 27

Subcase 2.2 Half-Venn diagrams of type B

To analyze this case we use both the half-Venn diagram and the corresponding half-Venn graph.

Corresponding to C_1 there are twelve vertices of $D(\mathfrak{F})$ from the matching M_1, so there will be four more in C_1. But by Remark 2, ${}^{I}C_1 \cap {}^{I}C_4 \cap {}^{I}C_5$ must contain four vertices, so there are no others inside C_1. Similar reasoning shows that there must be four vertices in each of the regions ${}^{I}C_1 \cap {}^{I}C_4 \cap {}^{O}C_5$, ${}^{I}C_1 \cap {}^{O}C_4 \cap {}^{I}C_5$, and ${}^{I}C_1 \cap {}^{O}C_4 \cap {}^{O}C_5$. This means that the vertices of $D(\mathfrak{F})$ must be distributed in one of the two ways shown in Figure 28, which we treat as subcases.

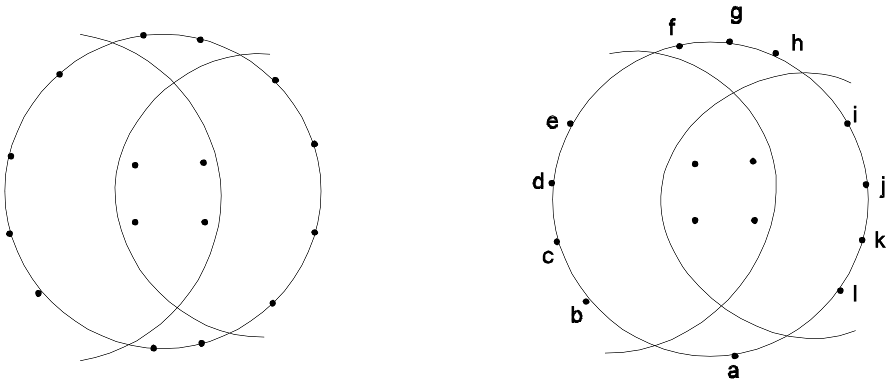

(a) **(b)**

Figure 28

Subcase 2.2.1 Vertices distributed as in Figure 28 (a)

The constraints imposed by Remark 2, of having only four vertices on each side and four vertices in the center region, force C_2 and C_3 to intersect in the center in one of two ways (up to symmetry), with additional segments as shown in Figure 29.

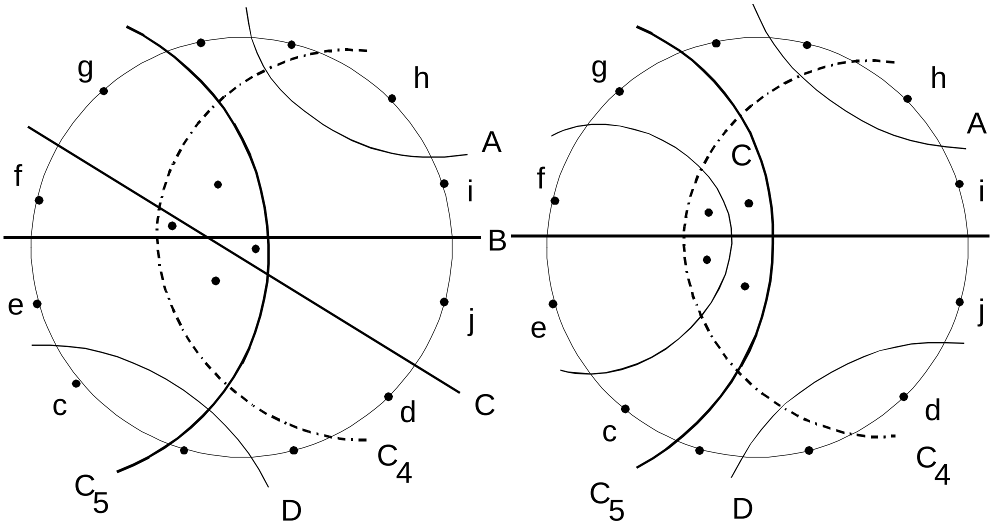

Figure 29

Remarks 1 and 3 show that segments A, C, and D must belong to the same curve, yielding a curve having six intersections with C_1, in contradiction to the assumption of Case 2. So Subcase 2.2.1 yields no Venn diagram.

Subcase 2.2.2 Vertices distributed as in Figure 28 (b)

The vertices b, a, and l must lie on a 4-cycle in the half-Venn graph, so b and l must both be adjacent to a vertex, say x, in the center. To complete the analysis, there are several vertices that could be chosen as key vertices; we choose f (symmetric with h). Remarks 1, 3, and 4 imply that the only vertex f could be adjacent with is k, which thus yields two subcases, according to whether f is adjacent with k or not.

Subcase 2.2.2(a) Vertex f is adjacent with k

The edge {g,j} is immediately forced, resulting in the situation in Figure 30 (a).
Then k, l, and x must lie on a 4-cycle, which is forced to contain a vertex in the center, say y. But now, vertices f, k, and y force edge {e,y}, yielding the configuration in Figure 30 (b). In this configuration, b and x must lie on a 4-cycle, but neither x nor y can have any more edges leaving the center; so x must be adjacent to another vertex, say v, in the center. We quickly deduce that the only possible configuration is that shown in Figure 30 (c).

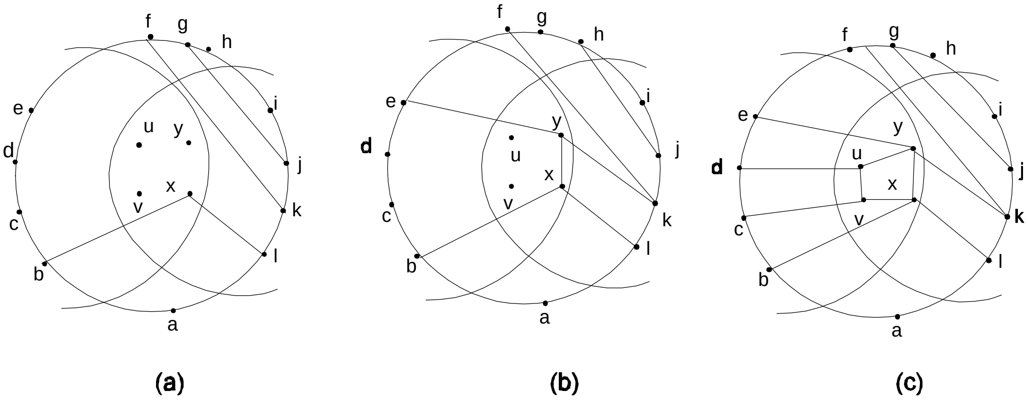

(a) (b) (c)

Figure 30

Subcase 2.2.2(b) Vertex f has no additional adjacencies

By symmetry, h has no additional adjacencies. Then g must be adjacent with both d and j. Now b (and l) can have no other adjacencies, so c, b, and x must lie on a 4-cycle with a vertex, say v, in the center, which by symmetry yields the configuration in Figure 31 (a). Then vertices v, c, and d must lie on a 4-cycle, which, together with the usual constraints quickly results in the unique configuration of Figure 31 (b).

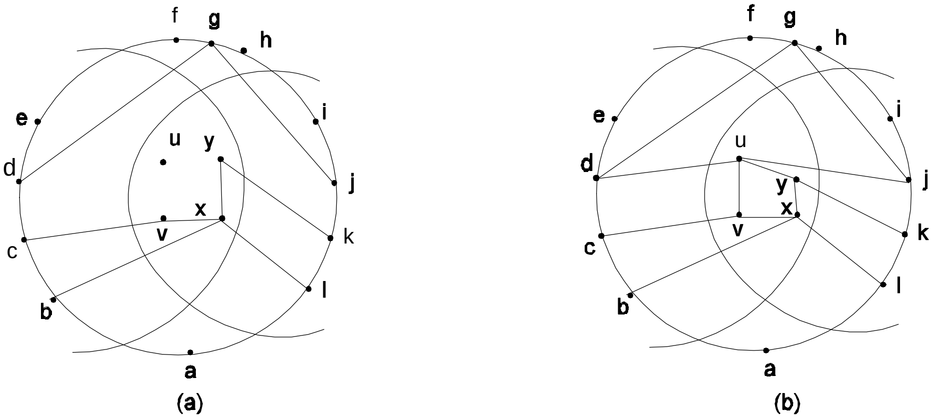

(a) (b)

Figure 31

Thus, case 2.2.2 yields the two half-Venn graphs shown in Figure 32 (a) and Figure 32 (b).

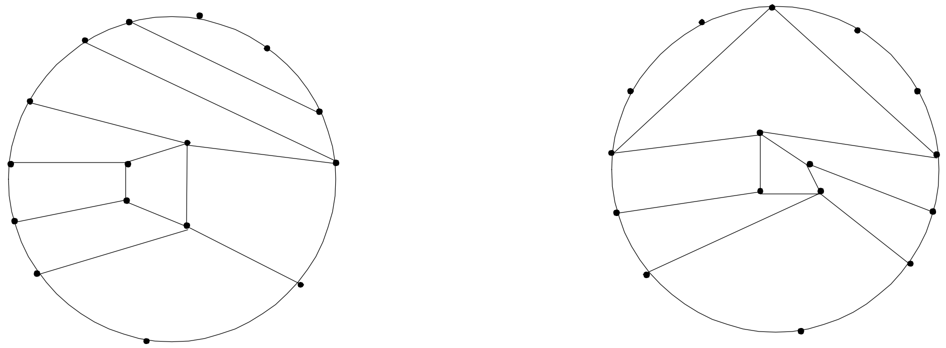

(a) (b)

Figure 32

We now proceed to the construction of the Venn diagrams resulting from this case. The two half-Venn graphs give the half-Venn diagrams in Figure 33.

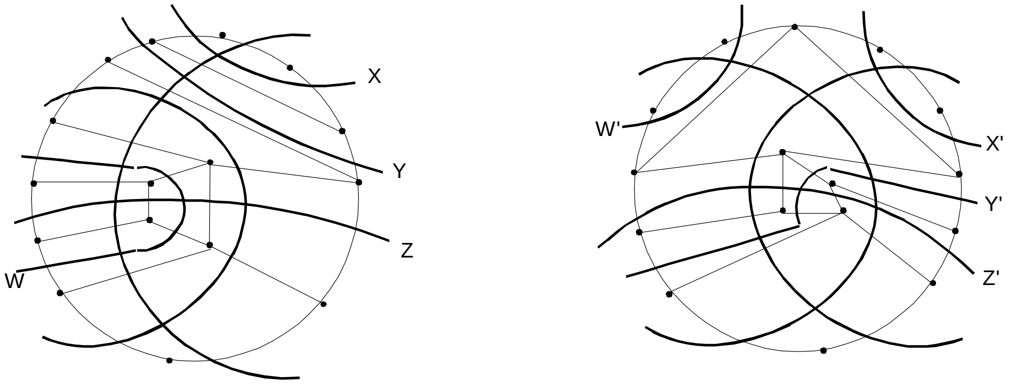

(a) (b)

Figure 33

The intersection numbers from Table 2 show that the curves C_4 and C_5 from one half-Venn diagram must each connect with one of C_4 or C_5 from another (possibly a copy of the same half-Venn). The fact that there is only one vertex at the bottom and three at the top of each half-Venn diagram requires that they be matched top-to-top, although they may be reflected left-to-right. The segments X, X', and W' form small closed curves if configuration a is matched with b or its reflection, if b is matched with itself or its reflection, or if a is matched with itself, so the only possibility is that a be matched with its reflection. This does in fact yield a Venn diagram as shown in Figure 34.

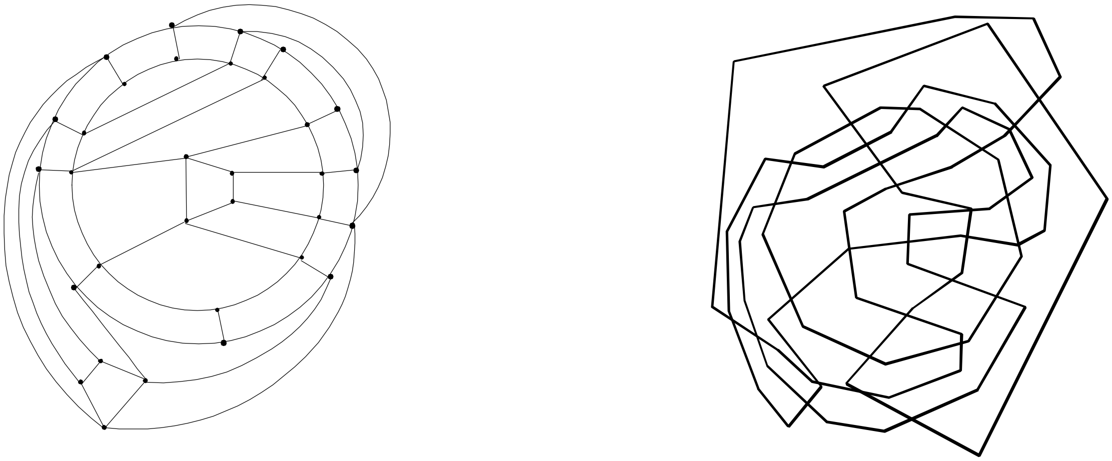

Figure 34

We summarize the results of Section 4.

PROPOSITION:

There is only one spherical Venn diagram having maximum intersection number 12.

This completes the proof of part i) of the theorem.

The analysis of simple, reducible 5-Venn diagrams with five curves was done in [6]; thus we need only to show parts (ii), (iii), and (iv) of the Theorem for irreducible diagrams.

Each diagram in Figures 1 and 2 is drawn as an exposed diagram proving the statement of (ii).

The unmarked diagram of Figure 1, and all of the diagrams of Figure 2 are drawn as convex diagrams proving that there are at least 11 convex, simple 5-Venn diagrams.

If in a Venn diagram the interiors of two curves intersect in a set with two (or more) connected components we will say that the Venn diagram has a *Grünbaum configuration*. An example is

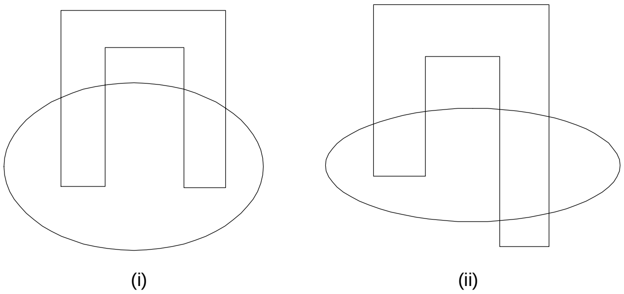

Figure 35

shown in Figure 35 (i). A *forbidden configuration* in a Venn diagram is two curves such that all the four possible intersections of the interiors and exteriors form a Grünbaum configuration. It is easy to see that if a Venn diagram has a forbidden configuration then any graph-isomorphic copy of the Venn diagram has a Grünbaum configuration, and thus there is no graph-isomorphic copy having a convex drawing in the plane. An example of a forbidden configuration is shown in Figure 35 (ii).

It is easy to find in diagram (iv) of Figure 1 a curve that forms a forbidden configuration with the dotted curve; thus it cannot be in the same class as any convex diagram. Any drawing of the diagrams (i, ii, iii) of Figure 1 with as an exposed diagram results in a drawing in which the dotted curve and another curve form a Grünbaum configuration. (We leave to check this for the reader.) Thus diagrams (i, ii, iii) of Figure 1 cannot be in the same class as a convex diagram. This completes the proof of (iii) of the Theorem.

To prove (iv) of the Theorem, we first need to recall the fact that a convex drawing of a Venn diagram in the plane must be exposed [4].

The unmarked diagram of Figure 1, all diagrams of Figures 2, and in [6] give 17 different convex, simple 5-Venn diagrams in the plane. To see that there are no more convex, simple 5-Venn diagrams in the plane we observe that any exposed drawings of a diagram from Figure 2 other than the given one results in a nonconvex drawing (having a Grünbaum configuration). We leave this for the reader to check.

To prove the last statement of (iv), we first observe that in Figures 2 and 3 all possible nonsimple convex irreducible 5-Venn diagrams have a drawing with five congruent ellipses. Since no other drawing of these diagrams in the plane is convex they cannot have a drawing with (congruent) ellipses.

This completes the proof of the Theorem. ▌

References

[1] K. B. Chilakamarri, P. Hamburger, and R. E. Pippert, *Venn Diagrams and Planar Graphs*, **Geometriae Dedicata 62** (1996), 73-91

[2] K. B. Chilakamarri, P. Hamburger, and R. E. Pippert, *Venn Diagrams: Announcement of Some New Results*, **Geombinatorics IV.4** (1995), 129-137

[3] D. B. West, *Introduction to Graph Theory*, **Prentice-Hall International**, **Inc**., (1996)

[4] B. Grünbaum, *Venn diagrams I*, **Geombinatorics I.4** (1992), 5-12

[5] B. Grünbaum, *Venn Diagrams and Independent Families of Sets*, **Mathematics Magazine**

1 (1975), 12-23

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

[7] P. Hamburger, *Planar Graph Models in Geometry, Combinatorics, Logic and Undergraduate Teaching*, **Mathematical Modeling in Undergraduate Curriculum**, **Proceedings of 1996 Second Symposium at the University of Wisconsin-La Crosse** (1997), 149-180

[8] P. Hamburger, *Constructing Venn Diagrams Using Graphs*, **Matematikai Lapok 4** (1994), 1-70, appeared in 1998, (in Hungarian with English Summary)

[9] P. Hamburger, *A Graph-Theoretic Approach to Geometry*, **Mathematics Magazine**, submitted for publication