

“11-set Doily.” Peter Hamburger’s 11-set “doily.” Top, the entire rotationally symmetric Venn diagram. Bottom, one of the 11 individual curves that makes up the diagram. (The other 10 curves are all, of course, rotated versions of this one.) (The mathematical foundation that made it possible to create the figures by artist Edit Hepp was made by Peter Hamburger; they are his intellectual properties; and he holds all the copyrights for this mathematics. The figures were created by artist Edit Hepp; they are her intellectual properties; and she holds all the copyrights.)

Combinatoricists Solve a Venn-erable Problem

Barry Cipra

VENN DIAGRAMS SEEM SO SIMPLE. Indeed, the topic is usually introduced in high school algebra, where overlapping circles are used to illustrate the various ways that sets can intersect (see Figure 1). Yet the subject offers a surprising number of mathematical challenges. One of these, regarding the existence of “rotationally symmetric” Venn diagrams, baffled mathematicians for over two decades, until it was recently solved—thanks in part to an inspired idea of an undergraduate math major. The detailed study of Venn diagrams falls under the heading of combinatorial geometry. Geometry because the diagrams consist of geometric shapes in the plane; combinatorial because they involve combinations of objects. Problems in combinatorics are often concerned with placing items in a list; in combinatorial geometry the problems are concerned with arranging items in space.

It helps to be precise about what a Venn diagram is. A Venn diagram for n sets consists of n closed curves. The “curves” need not be smooth—they may in fact be polygonal—and they need not all have the same size or shape (see Figure 2, next page). Each pair of curves can (and indeed must) intersect at one or more points, but not infinitely often. In particular, there are no “shared arcs” in a Venn diagram. When all drawn at once, the curves cut the plane into regions that are *inside* certain curves and *outside* the others. The crucial, defining feature of a Venn diagram is that each such inside/outside combination is represented by exactly one such region.

The number of distinct regions in a Venn diagram grows exponentially: With n sets, there are 2^n regions, including the common interior and common exterior. This is where the complications begin to creep in: As n gets larger and larger, those 2^n regions, it turns out, can be arranged in lots and lots of different ways.

This is not at all evident in the familiar 2- and 3-set pictures. Indeed, there is essentially only one way to draw a Venn diagram with two sets: No matter how wildly you draw the two curves, they can be thought of as just a badly drawn pair of circles. (This is all the more true of a 1-set Venn diagram—and totally trivial if there are no sets!) However, there are 14 differ-




Figure 1. The familiar 2-set and 3-set Venn diagrams using circles. A Venn diagram consists of n simple closed curves that divide the plane into 2^n connected regions, in such a way that each “inside-outside” combination is represented exactly once. For instance, if there are two curves A and B , the four combinations are “inside-inside,” “inside A -outside B ,” “outside A -inside B ,” and “outside-outside.” Both of these diagrams are rotationally symmetric. (Figure from “Venn Diagrams and Symmetric Chain Decompositions in the Boolean Lattice,” J. Griggs, C. E. Killian, and C. D. Savage, *Electronic Journal of Combinatorics* 11 (2004), Research Paper 2, 30pp. electronic.)

ent ways to draw a 3-set Venn diagram, grouped into 6 classes (see Box, “Vennis Anyone” and Figure 3, page 44).

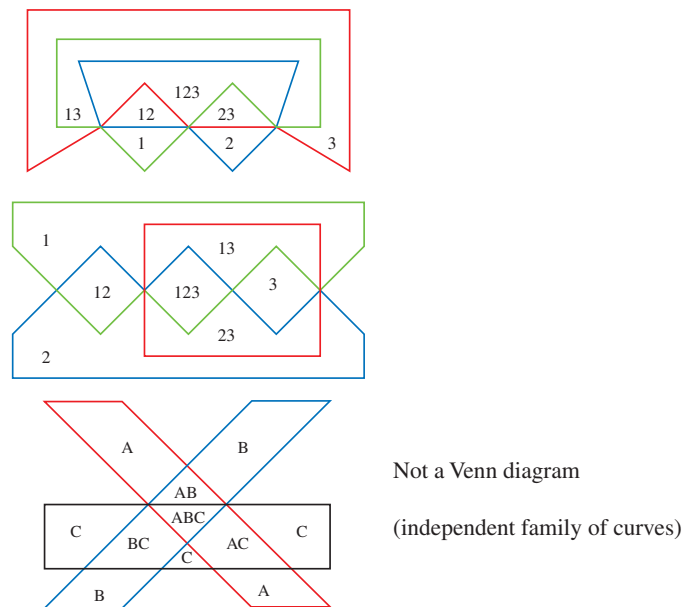


Figure 2. Venn diagrams do not have to be drawn with smooth curves. Here are two polygonal 3-set Venn diagrams, and one “non-Venn diagram.” In the non-Venn diagram there is an inside-outside combination that consists of two disjoint regions; can you find it? (Figure from “Venn Diagrams and Symmetric Chain Decompositions in the Boolean Lattice,” J. Griggs, C. E. Killian, and C. D. Savage, *Electronic Journal of Combinatorics* 11 (2004), Research Paper 2, 30pp. electronic.)

When there are more than three sets, the enumeration problem is unsolved, except for a special type of diagram known as a *simple* Venn diagram. A Venn diagram is said to be simple if every point of intersection has exactly two curves crossing. A famous theorem known as Euler’s formula, applied to simple Venn diagrams, shows that there are $2^n - 2$ curve crossings in a simple n -set Venn diagram (see Box, “Venn Meets Euler,” p. 45). For three sets, there is just one simple Venn diagram, as Figure 3 shows. For four sets, there are two simple Venn diagrams, both belonging to the same class. For $n = 5$, the number of classes of simple Venn diagrams jumps to 19. Beyond five, the number of classes has not been tallied.

In 1963, David Henderson at Swarthmore College took note of another property that is obviously true for the familiar 2- and 3-set Venn diagrams, but less obvious—and, indeed, *not* true—in general: The familiar diagrams are “rotationally symmetric.” That is, there is a point in the region corresponding to the common intersection ($A \cap B$ for $n = 2$, and $A \cap B \cap C$ for $n = 3$) about which the diagram is symmetric under rotation by $360/n$ degrees (or $2\pi/n$ radians). More precisely, a Venn diagram is rotationally symmetric if it can be drawn by drawing one close curve and then rotating it by multiples of $2\pi/n$ radians about a

point in its interior to produce the other curves. Are there such diagrams, Henderson asked, for values of n greater than three?

Vennis, Anyone?

Venn diagrams are traditionally drawn on paper. Maybe they should really be drawn on tennis balls.

The plane and the sphere are closely connected. In particular, the sphere can be thought of as the plane with an extra “point at infinity.” (Alternatively, the plane is a “punctured” sphere.) The connection comes courtesy of a correspondence called stereographic projection: Place a transparent globe so that its south pole sits on an (infinite) tabletop, and shine a laser pointer from its north pole to a point on the table. (See “Stereographic Projection,” p. 51.) On its way, the laser beam passes through a point on the globe. That is the unique point on the sphere that corresponds to the chosen point on the table, and vice versa. As the point on the table moves further and further away from where the globe is sitting, its corresponding point on the sphere moves closer and closer to the north pole. Thus the north pole itself corresponds to the “point at infinity” for the plane.

If you draw a Venn diagram on a sphere, say using magic markers on a tennis ball, the curves separate the sphere into 2^n regions (n being the number of curves). But there is no longer an obvious “inside” and “outside” for each curve; instead, there is just “the side that contains the north pole” and “the side that doesn’t.” In a planar Venn diagram, there is one region that appears much different from the others—the unbounded region, which surrounds the whole diagram. But on a sphere, there is no such distinction; the region that contains the north pole looks pretty much like all the others.

This is actually an advantage because one is free to take the sphere and reposition it so that some other point, inside one of the other regions, becomes the north pole, and then re-project onto the plane. The result is another Venn diagram. It may, of course, be equivalent to the first, but it may not. When $n = 4$, for example, there are regions bounded by three curves and regions bounded by four, so there are two distinct projections, depending on which type of region is chosen to contain the north pole. The various diagrams, while “planely” different, are said to belong to the same class. That is why, in the enumeration problem, the number of *classes* of Venn diagrams with n sets is smaller than the absolute number of Venn diagrams.

The answer—or part of it—is “not if n is composite.” The reason, as Henderson showed, is fairly simple. Aside from the common interior and common exterior of the n curves, each region rotates through a total of n copies of itself, each of the same “type,” meaning the number of sets whose intersection is that region. This means that n must divide the number of pairwise intersections, the number of three-way intersections, etc., on up to the number of $n - 1$ -way intersections. These numbers are well known as the binomial coefficients $\binom{n}{k} = n!/k!(n - k)!$

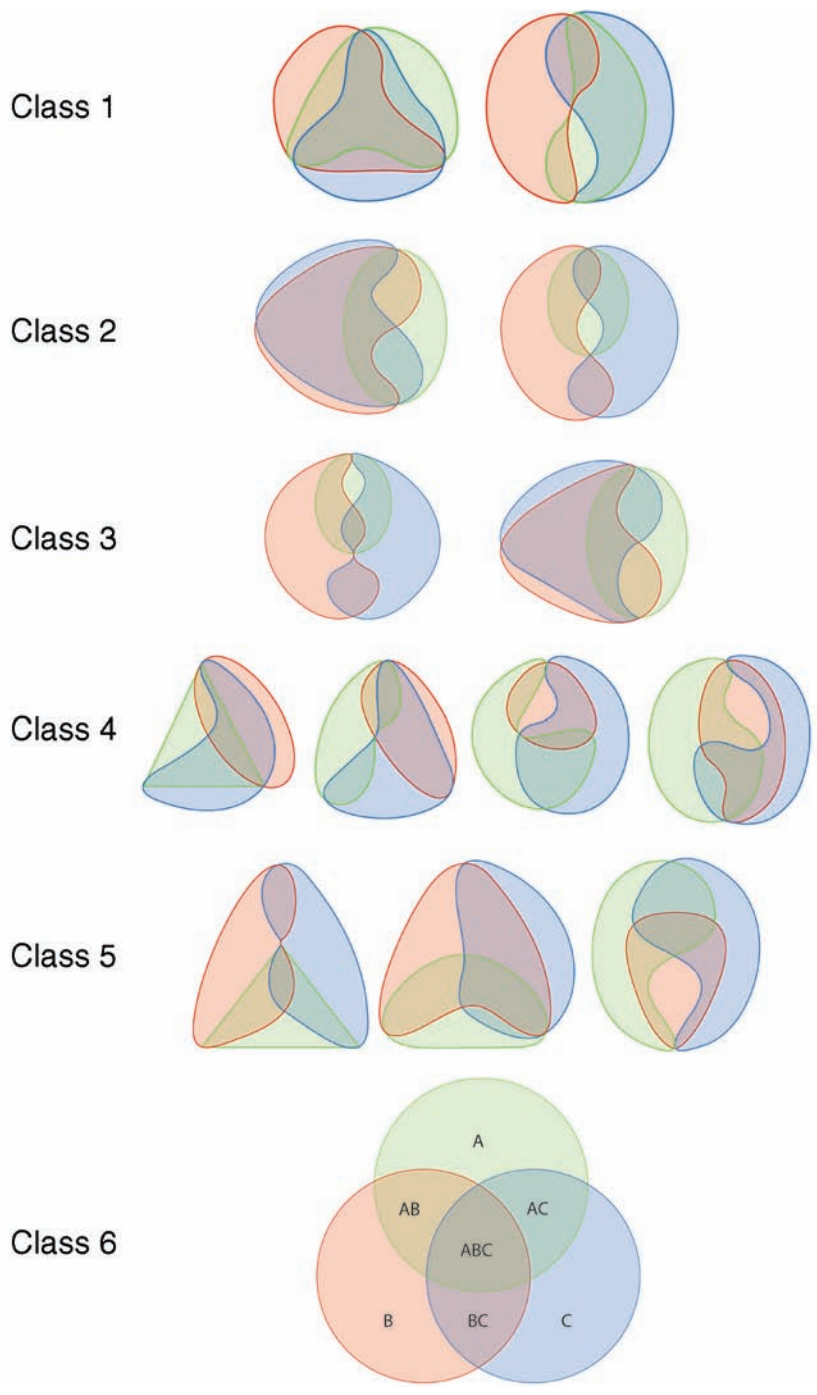


Figure 3. The 14 different 3-set diagrams. A “simple” Venn diagram has no threefold crossings; note that only one of the 14 3-set diagrams is simple. (Based on original figures supplied by Frank Ruskey at the University of Victoria.)

for $k = 2, 3$, up to $n - 1$ (see Figure 4). But a famous theorem, originally proved by Leibniz, says that n divides all its binomial coefficients if and only if n is prime. For example, $n = 4$ does not divide $\binom{4}{2} = 6$. So the six regions that correspond to pairwise intersections cannot be arranged around a central point with fourfold symmetry.

Venn Meets Euler

The great Swiss mathematician Leonhard Euler was actually one of the first mathematicians to study Venn diagrams—more than 100 years before John Venn took up the subject! Euler also studied more general types of *planar graphs*, which consist of curves, called edges; connecting points, called vertices; separating the plane into distinct regions, called faces. (A graph is called planar when none of its edges cross.) Euler showed that, for any planar graph, the numbers of vertices V , edges E , and faces F satisfy the formula $V - E + F = 2$.

For Venn diagrams, the number of faces is 2^n . For simple Venn diagrams, it is easy to see that there are twice as many edges as vertices, i.e., $E = 2V$ (see Figure 1, p. 41). Thus Euler's formula implies $V - 2V + 2^n = 2$, or $V = 2^n - 2$.

Leibniz's theorem leaves open the possibility that prime values of n do permit rotationally symmetric Venn diagrams. Henderson gave two examples for $n = 5$, one with irregular pentagons and one with quadrilaterals. He also claimed to have an example for $n = 7$ with irregular hexagons, but was unable to reproduce it later.

Branko Grünbaum at the University of Washington took up the problem in 1975. He gave additional examples for $n = 5$, including one using ellipses (see Figure 5). Frank Ruskey at the University of Victoria later did an exhaustive computer search for symmetric Venn diagrams with five sets. He found there are 243 different examples, of which only one—Grünbaum's ellipses—is simple.

Grünbaum returned to the problem in 1992, when he discovered a 7-set example (see Figure 6, page 46). Shortly after, additional 7-set diagrams were found by Anthony Edwards at Cambridge University and, independently, by Peter Winkler at Dartmouth College and Carla Savage at North Carolina State University. With these examples, it became natural to conjecture that rotationally symmetric Venn diagrams with n sets exist for *all* primes n . But, there, things came to an apparent impasse: Venn diagrams with 5 or 7 sets are still small enough that a combination of cleverness and patience suffices to ferret out examples of symmetry, but with 11 sets the number of possibilities to pick through becomes intractably large. As Grünbaum put it in a status report in 1999, "The sheer size of the problem for 11 curves puts it beyond the reach of the available approaches through exhaustive computer searches."

Enter Peter Hamburger. A graph theorist at the University of Indiana-Purdue University at Fort Wayne, Hamburger had worked on some of the enumeration problems for Venn diagrams. (In a series of papers with Raymond Pippert and Kiran




Figure 4. Pascal's triangle can be used to prove that there are no rotationally symmetric Venn diagrams with 4 sets. The k -th entry in the n -th row of Pascal's triangle (counting the 1 at the top as the zero-th row) enumerates how many regions are inside k curves but outside $(n - k)$ curves. If the Venn diagram is rotationally symmetric, each of the entries (not including the 1s at the beginning and end) must be divisible by n . But the middle entry of the fourth row is not divisible by 4. In fact, for any non-prime n , the n th row of Pascal's triangle contains some elements not divisible by n , hence a Venn diagram with n -fold rotational symmetry is impossible.




Figure 5. Branko Grünbaum's rotationally symmetric Venn diagram with 5 regions, all bounded by ellipses. (Figure courtesy of Branko Grünbaum.)




Figure 6. Grünbaum's rotationally symmetric Venn diagram with 7 regions. (Figure courtesy of Branko Grünbaum.)

Chilakamari, he identified the 14 different Venn diagrams with three sets and the 19 classes of simple Venn diagrams with five sets.) Hamburger thought there might be a way to construct rotationally symmetric diagrams using methods developed for a different, purely combinatorial problem: symmetric chain decompositions of Boolean lattices.

A Boolean lattice on n elements can be viewed as the subsets of the set $\{1, 2, \dots, n\}$, stratified by the size of the subsets, and with each subset "connected" to the subsets one level above and below it that differ by the additional presence or absence of one element (see Figure 7a). A "chain" is simply a string of connected subsets, and a "chain decomposition" is a collection of disjoint chains that account for all the subsets. (Some of the "chains" may consist of a single subset. In fact, letting each subset constitute its own chain is a perfectly good chain decom-

position.) A chain decomposition is “symmetric” if the smallest and largest subset in each chain are of complementary size—that is, if the smallest subset in a chain has k elements, then the largest subset in that chain has $n - k$ elements (see Figure 7b).




Figure 7. A chain decomposition of a Boolean lattice is a partition into several “chains,” in which each chain consists of a collection of sets that differ by turning one “inside” into and “outside” or vice versa. It is trivial to draw chain decompositions if you are given a Venn diagram. The key step in constructing rotationally symmetric Venn diagrams is to do the reverse: to go from an abstract chain decomposition with the right kind of symmetry properties to constructing a concrete Venn diagram. Peter Hamburger showed how to do this for an 11-set Venn diagram, and Carla Savage, Peter Griggs, and Chip Killian accomplished it for all remaining primes. (Figure from “Venn Diagrams and Symmetric Chain Decompositions in the Boolean Lattice,” J. Griggs, C. E. Killian, and C. D. Savage, *Electronic Journal of Combinatorics* **11** (2004), Research Paper 2, 30pp. electronic.)

Starting with a Venn diagram, it is easy to draw chains by connecting adjacent regions in the diagram (always crossing curves from inside to out), and therefore easy to create chain decompositions. If the diagram is rotationally symmetric, the

Hamburger was able to construct an 11-set “Doily”—his term for the kind of Venn diagram that comes from his method, because of their lacy look—that is rotationally symmetric.

chains (ignoring the full and empty sets) can be chosen so as to rotate onto one another. The result is not guaranteed to be a symmetric chain decomposition. However, sometimes it is.

Hamburger’s idea was to go the other way: Start with a symmetric chain decomposition and systematically build a Venn diagram around it. Moreover, by imposing an additional rotational symmetry on the chains, the resulting diagram would be rotationally symmetric.

There was no guarantee this approach would work, but the two symmetry conditions and the systematic construction sharply reduced the number of possibilities to be considered. Examples for $n = 5$ and 7 turn out to be easy to find. Hamburger then set his sights on the next, unsolved case: $n = 11$.

The method worked. Hamburger was able to construct an 11-set “Doily”—his term for the kind of Venn diagram that comes from his method, because of their lacy look—that is rotationally symmetric (see “11-set Doily,” p. 40). His wife, Edit Hepp, who is an artist, has turned several of his examples into works of art. Hepp uses colored pencils to fill in regions corresponding to different types of intersections. The results are beautiful, mandala-like abstracts.

Hamburger’s approach opened a new route for constructing rotationally symmetric Venn diagrams, but it still required cleverness and patience in picking through possibilities—enough of them that the next case, $n = 13$, was too dense a thicket for Hamburger to penetrate. It ultimately took an undergraduate to find a way.

Jerry Griggs at the University of South Carolina read Hamburger’s paper and started thinking about how to obtain Venn diagrams from symmetric chain decompositions. He found that the construction depended on what he called a chain-covering property, which he showed could always be satisfied, even for composite n . To get rotational symmetry for prime values of n , it would be necessary (and sufficient) to impose rotational symmetry on the chain decomposition while maintaining the chain-covering property.

Imposing symmetry amounts to deciding which regions occur in a single “wedge” of the diagram and then requiring the next wedge contain the same regions “moved up” by one. For example, if one wedge contains the region corresponding to the interiors of curves 1, 3, and 6 (and the exteriors of curves 2, 4, 5, and 7, for $n = 7$), then the next wedge must contain the region corresponding to the interiors of curves 2, 4, and 7, the wedge after that must contain the region corresponding to the interiors of curves 3, 5, and 1, and so forth.

A nice way to abbreviate all this is to use binary strings to describe the regions: 1010010 is in one wedge, 0101001 in the next, 1010100 in the one after that, and so forth. When n is prime, every such string, except for all 1’s or all 0’s (corresponding to the center and the unbounded, outer region of the Venn diagram), belongs to an “orbit” of n strings. For a diagram to be symmetric, each wedge must include exactly one “representative” of each orbit. Consequently, Griggs saw, the key to rotational symmetry was to find a rule for picking these representatives in a way that preserved the crucial chain-covering property.

Like Hamburger, Griggs got stuck. “None of the rules I tried for picking representatives worked out,” he says. There was another problem as well, he adds: “I never had much time to work on it.”

Griggs described his ideas to Carla Savage at North Carolina State University, and Savage enlisted an undergraduate, Charles “Chip” Killian (currently a graduate student at Duke University), to work on the next case, $n = 13$. “I thought we could hack out 13,” she says. It seemed like a good project for a student.

Killian hacked out a lot more than 13. He found a rule for picking representatives that worked for *all* primes.

Killian’s rule is surprisingly simple. It is based on what is called the “block” structure of the binary string abbreviations for the regions. For example, to determine which string in the cycle generated by 01110011010 (for $n = 11$) belongs to the wedge whose outermost bound is an arc of curve number one, consider the six rotations of it that begin with a 1: 11100110100, 11001101001, 10011010011, 11010011100, 10100111001, and 10011100110. (Everything in this wedge is in the interior of the first curve, which is why it’s not necessary to consider strings beginning with a 0.) In each of these candidate strings there are alternating blocks of 1’s and 0’s. Killian’s rule is to pick the one with the least total number of 1’s and 0’s in the first pair of blocks—in this case, 10100111001, for which the total is 2. If there is a tie, the rule is to look at the next pair of blocks, and so forth, until the tie is broken, which is guaranteed to happen eventually.

When Killian first showed Savage his rule, she was skeptical. “It took a while for him to convince me we should take this seriously,” Savage recalls. It isn’t obvious the rule is copacetic with the chain-covering property. (It’s also not obvious the tie-breaking procedure always works, though that’s fairly easy to prove.) But within a few months, Griggs, Savage, and Killian had checked that everything worked. “There are a lot of things to verify, but once you see them they’re easy to prove,” Savage says.

It’s somehow fitting that the last step in solving Henderson’s question about rotationally symmetric Venn diagrams should be taken by a student: Henderson wrote his paper when *he* was an undergraduate. However, Hamburger’s doilies and the general solution by Griggs, Savage, and Killian do not settle all the questions that can be asked about these diagrams. In particular, the constructions based on symmetric chain decompositions are highly non-simple: There are points, in fact, where all n curves cross. (These points define the wedges.) In general, the constructions based on symmetric chain decompositions give rise to diagrams with $\binom{n}{(n-1)/2}$ vertices, which is much smaller than the $2^n - 2$ vertices in a simple Venn diagram for n curves.

Hamburger and colleagues György Petruska at Indiana-Purdue University Fort Wayne and Attila Sali at the Alfréd Rényi Institute of Mathematics in Budapest have found ways to tease apart many of the multi-curve crossings in the $n = 11$ case. Their best result to date has 1837 vertices—167 in each wedge—or 209 shy of the 2046 vertices an 11-curve Venn diagram needs to be simple. (Hamburger’s original 11-set Doily

It’s somehow fitting that the last step in solving Henderson’s question about rotationally symmetric Venn diagrams should be taken by a student: Henderson wrote his paper when *he* was an undergraduate.




Figure 8. Grünbaums Venn diagram with five equilateral triangles. (Figure courtesy of Branko Grünbaum.)


has $\binom{11}{5} = 462$ vertices.) Together with Frank Ruskey and Mark Weston at the University of Victoria, Savage and Killian have shown it's possible to increase the number of crossings in their general construction to at least 2^{n-1} , but that's still barely halfway to the $2^n - 2$ crossings for a simple Venn diagram. To date there are no known examples of a simple rotationally symmetric Venn diagram when n is greater than 7.

Venn diagrams offer many other challenges. In 1984, Peter Winkler, who is now at Dartmouth College, conjectured that every simple Venn diagram with n curves can be extended to a simple Venn diagram with $n + 1$ curves by the addition of one more curve. (Grünbaum noted that extendability had not been proved even if simplicity were not assumed. This was settled in 1996 by Hamburger, Pippert, and Chilakamarri, who showed that every Venn diagram with n curves can be extended to one with $n + 1$ curves. Their proof, however, produces non-simple extensions.)

Also in 1984, Grünbaum asked whether it is possible to draw a Venn diagram with six triangles. Grünbaum and Winkler had solved the corresponding problem for $n = 5$. There is, in fact, a simple, rotationally symmetric Venn diagram with five equilateral triangles (see Figure 8). This was settled in 1999 by Jeremy Carroll, a research scientist at Hewlett-Packard Laboratories in Bristol, England. By means of an exhaustive -computer search, Carroll found there are 126 different Venn diagrams with six triangles. His method, however, does not answer a related question: Can any of these diagrams be drawn using six equilateral triangles? Problems like these are likely to keep researchers—and their students—busy for years to come.

A Venn Diagram Roundup

Venn diagrams have so many interesting properties—in addition to rotational symmetry, simplicity, and extendability, there is monotonicity, convexity, rigidity, reducibility, and on and on—it would take a complicated, multi-curve Venn diagram to keep track of them all. A good place to start for all your Venn diagram needs is the online “Survey of Venn Diagrams,” by Frank Ruskey at the University of Victoria. Originally published in 1997 in the *Electronic Journal of Combinatorics* (which is also where the Griggs-Savage-Killian paper appears), Ruskey’s survey (www.combinatorics.org/Surveys/ds5/VennEJC.html) has been updated to include many of the latest results.



Stereographic Projection. This figure illustrates a standard method for building a correspondence between the points on the sphere (minus the point at the north pole) and all the points on the plane. The three points labeled A , C and E on the sphere are projected onto the points B , D and F , respectively, on the (infinite) table. Likewise, the circle passing through A , C , and E projects down to the circle through B , D , and F . If the latter circle were part of a Venn diagram drawn on the table, it could be lifted back up to become part of a corresponding Venn diagram drawn on the sphere.