From *SIAM News*, Volume 37, Number 1, January/February 2004

Venn Meets Boole in Symmetric Proof

By Barry A. Cipra

Venn diagrams, long a staple of high school algebra, have also become a staple—or at least a paperclip—for combinatorial geometers. The simplicity of the familiar two- and three-circle Venn diagrams turns out to be deceptive. Once a fourth set is added, circles no longer suffice and the subject raises a raft of challenging problems. One of them—the existence or non-existence of rotationally symmetric Venn diagrams—was solved only recently.

Jerry Griggs of the University of South Carolina (who a year ago became the editor-in-chief of *SIAM Journal on Discrete Mathematics*), Carla Savage of North Carolina State University, and her student Charles "Chip" Killian (now a graduate student at Duke University) have shown that it's possible to produce rotationally symmetric Venn diagrams whenever the number of sets is prime. Their proof builds on a breakthrough by Peter Hamburger of Indiana University– Purdue University Fort Wayne. Savage described the result last January at the workshop Algorithms for Listing, Counting, and Enumeration (ALICE03), which was held in conjunction with the SIAM/Association for Computing Machinery Symposium on Discrete Algorithms.

The Carolina researchers' proof for primes polishes off the problem, because it was already known that rotationally symmetric Venn diagrams cannot be drawn when the number of sets is composite. This part of the proof is easy, but understanding it requires a proper explanation of just what a Venn diagram is (and what it means to be rotationally symmetric).

In a Venn diagram for n sets, each set is represented as the interior of a simple closed curve. A simple curve doesn't cross itself, as in a figure eight—that is, the interior is connected. Connectedness is the crucial property for the intersections as well: For each choice of k sets, the region inside the chosen sets (and outside the others) must be connected. It must also be non-empty. (Diagrams in which some intersections fail to occur are called Euler diagrams. They first appeared in Euler's famous *Letters to a German Princess*.)

A Venn diagram is rotationally symmetric (Venn diagrammers usually omit the adverb and just say "symmetric") when all the curves are obtained by rotating one of them about a common center by multiples of $2\pi/n$ radians. Two- and three-circle Venn diagrams can obviously be represented as rotationally symmetric Venn diagrams. To create a symmetric Venn diagram for n sets, all you have to do, it would seem, is draw a curve, pick a center, rotate the curve to get the rest of the sets, and then check that you actually get all 2^n intersections (including the common exterior).

The problem is, you often don't get all the intersections. And when you do, they're often disconnected. For composite n, in fact, it's impossible to get a complete set of non-empty, connected intersections no matter what you do. The reason is simple: Connectedness implies precisely $\binom{n}{k}$ regions representing the in-tersections of k sets at a time, and rotational symmetry requires that n divide this number for k between 1 and $n - 1$, which, by a theorem of Leibniz, happens only when n is prime.

That leaves the problem of primes. The question was first raised by David Henderson of Swarthmore College in 1963. Henderson gave two polygon examples for $n = 5$, one with irregular pentagons and the other with quadrilaterals. He also claimed to have found an example for $n = 7$ using irregular hexagons, but was later unable to reproduce it.

In 1975, Branko Grunbaum of the University of Washington gave more examples for $n = 5$, including one with ellipses, but opined that there were no rotationally symmetric Venn diagrams for $n = 7$. The problem then languished until 1992, when Grunbaum, trying to prove his negative conjecture, succeeded instead in constructing an example.

Grunbaum's $n = 7$ construction persuaded most people that rotationally symmetric Venn diagrams exist for all prime n, but it also hinted that finding them would be prohibitively difficult. As Grunbaum wrote in a 1999 survey of the subject, "The sheer size of the problem for 11 curves puts it beyond the reach of the available approaches through exhaustive computer searches."

Enter Hamburger. In a paper titled "Doodles and Doilies, Non-Simple Symmetric Venn Diagrams," presented in 1999 at a conference and published in 2002 in *Discrete Mathematics*, Hamburger introduced a new approach to the problem that enabled him to solve the $n = 11$ case—with no computer search at all. Hamburger calls his rotationally symmetric diagrams "doilies" for the lacy crocheted pieces they resemble (see Figure 1). His "doodles" are the key. Roughly speaking, a doodle is a compact blueprint for a $2\pi/n$ wedge of the doily. Consisting of loops nested within loops, a doodle is derived from a study of

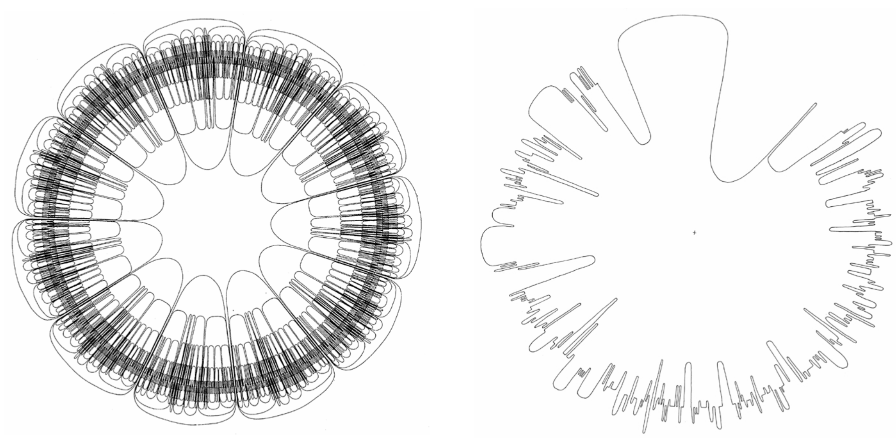

Figure 1. *The non-simple 11-doily (left). The simple closed curve and the center of rotation (right). The rotation of this curve 11 times over 360/11 degrees creates the non-simple 11-doily at left. Courtesy of Edit Hepp and Peter Hamburger.*

symmetric chain decompositions of the Boolean lattice. Indeed, Hamburger sees Boolean lattices as the main motivation for his work on Venn diagrams.

The Boolean lattice B_n is the set of all subsets of {1,2, ... ,n}, partially ordered by inclusion. It can also be thought of as the n-dimensional hypercube, with vertices running from all 0's to all 1's. A symmetric chain is a path in the hypercube whose starting point has as many 0's as its endpoint has 1's; each step along the way (in one direction) turns a 0 into a 1, which is equivalent to enlarging a subset in B_n by one item. A symmetric chain decomposition (SCD) separates the lattice into distinct chains. SCDs are useful objects in the theory and application of Boolean lattices (see Figure 2).

It's not too hard to turn a symmetric chain decomposition into a decomposition of the plane whose 2^n regions correspond to subsets of {1,2, ... ,n}. There are two tricky parts. One is to pick an SCD that produces a Venn diagram, rather than a jumble of disconnected duplicate intersections. The other is to get a diagram that is rotationally symmetric. But the real trick is to accomplish the two simultaneously. That's what Hamburger did for $n = 11$, and what Griggs et al. have accomplished in general.

After a warmup with $n = 7$, Hamburger found a decomposition of B_{11} (actually a subposet thereof) into symmetric chains whose doodle produced a doily. His construction pointed the way toward a general solution, but it was not clear that appropriate doodles would necessarily exist for larger prime numbers, or how to find them if they did. That's where Griggs came in.

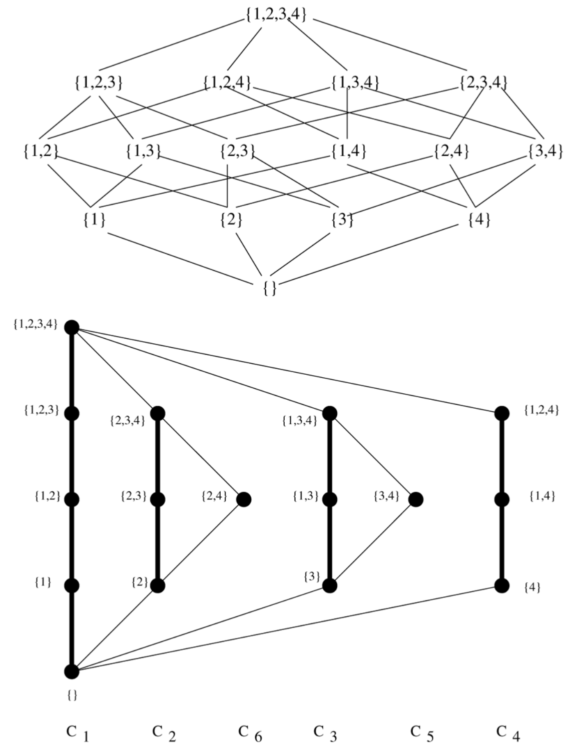

Figure 2. *The Boolean lattice B_4 (top) and a symmetric chain decomposition with the chain cover property (bottom). From Jerrold Griggs, Charles E. Killian, and Carla D. Savage, "Venn Diagrams and Symmetric Chain Decompositions in the Boolean Lattice."*

Although he mentioned being inspired by an SCD introduced by Curtis Greene of Haverford College and Daniel Kleitman of MIT, Hamburger did not make explicit use of their decomposition in his paper. Griggs "began playing around" with the Greene–Kleitman SCD for $n = 5$ and 7 and found that obtaining Venn diagrams depended on a "chain covering property" of the decomposition. The rotational symmetry, Griggs found, required picking representatives of the rotational classes (e.g., 11000, 01100, 00110, 00011, and 10001 are equivalent under rotation) in a way that maintained the chain cover property.

"I believed we only had to find a rule to pick the right representative from each orbit, so that they could be joined together by the [Greene–Kleitman] method to obtain an SCD with the chain cover property, yielding the desired Venn diagram for all primes n," Griggs says. "The trouble was, none of the rules I tried for picking representatives worked out, and I never had much time to work on it."

Enter Savage—and Killian.

In the mid 1990s, Savage had tried her hand at the $n = 11$ case. When Griggs told her about Hamburger's idea for using symmetric chain decompositions and his own idea of the chain covering property, "I thought we could hack out 13," she says. She also thought it would be a good problem for an undergraduate to work on.

Indeed, it was Killian who came up with the final breakthrough—the rule that had eluded Griggs for picking representatives from the rotational orbits. And not just for $n = 13$, but for *all* prime n.

"It took a while for him to convince me we should take this seriously," Savage recalls. But within a few months, the three researchers had checked all the details. "There are a lot of things to verify, but once you see them they're easy to prove," Savage says.

The "chain cover property" part of the proof actually works for all n, prime or not. The construction produces Venn diagrams that are "monotone"—meaning that each region corresponding to the intersection of k sets (for $0 < k < n$) is adjacent to regions corresponding to intersections of both $k - 1$ sets and $k + 1$ sets. By a 1998 result of Bette Bultena, Grunbaum, and Frank Ruskey, it also means that the diagram can be redrawn with convex curves. (Their result is actually an if-and-only-if result: A Venn diagram is monotone if and only if it can be drawn with convex curves. Hamburger's $n = 11$ construction is also monotone. But to redraw it with convex curves would make almost all the regions too small to see—the 11 curves would look like a thick, fuzzy, slightly undulating circle.)

In a monotone Venn diagram a lower bound for the absolute minimum number of vertices to be $\binom{n}{n/2}$ (e.g., 6 when $n = 4$) is established by a famous (1921) theorem of Sperner. In a 1998 paper, Bultena and Ruskey showed that this lower bound is sharp for all integer n. The Carolina team's construction achieves precisely this lower bound, with a simpler proof.

Meanwhile, Hamburger, in collaboration with Attila Sali of the Renyi Institute of Mathematics in Budapest, has used his method to generate 11-doilies with as few as 231 vertices (decidedly non-monotone) and as many as 1001. Hamburger's wife, Edit Hepp, who is an artist, has turned a number of them into mandala-like color prints. (Hamburger's Web site, http://www.ipfw.edu/math/hamburger/, has a page of pictures taken at an exhibition they put on at a local coffee house.) The two have written an article about their venture for the art–science journal *Leonardo*.

The "ideal" 11-doily would have 2046 vertices—that is, only two curves would intersect at each vertex. For $n = 3, 5$, and 7, such "simple" symmetric Venn diagrams are known. Indeed, Ruskey (whose survey paper in the *Electronic Journal of Combinatorics* (http://www.combinatorics.org/Surveys/ds5/VennEJC.html) is the place to go for information on Venn diagrams) has identified 56 such diagrams for $n = 7$, including all that are monotone. There are 23 of them. (The list of 33 non-monotone simple diagrams for $n = 7$ is thought to be complete, but that remains to be proved. There is only a partial list of non-simple examples. For $n = 5$ there is just one simple symmetric Venn diagram, and 242 that are non-simple.) But all of Hamburger's 11-doilies and all those generated by the Carolina crew's method are non-simple. For a seemingly trivial topic in high school algebra, Venn diagrams continue to pose a lot of unsolved problems.

Barry A. Cipra is a mathematician and writer based in Northfield, Minnesota.