Quite Easily Done

The line between easy mathematical problems and hard ones is finely drawn. Some problems seem to cross back and forth: First they look easy, then they seem hard, and then, when they're finally solved, they look easy again. A recent example is a simple-sounding combinatorial puzzler called the Dinitz problem. First posed in 1978, the Dinitz problem has finally been solved with a surprisingly simple proof, but only after fifteen years during which it seemed a very tough nut to crack.

The story starts in the late 1970s. Jeff Dinitz, then a graduate student at Ohio State University (now a professor at the University of Vermont), was studying properties of combinatorial arrangements known as latin squares. A latin square is an $n \times n$ array of n symbols—say a 5×5 array of stars, squares, circles, diamonds, and triangles—in which no symbol appears more than once in any row or column (see Figure 1). Latin squares are useful, for example, in the design of experiments, to protect against bias. If, say, you want to compare five different herbicides in a corn field, but want to make sure the results aren't affected by variations in soil quality from one side of the field to another, then dividing the field into a 5×5 latin square pattern is an efficient way to design the experiment.

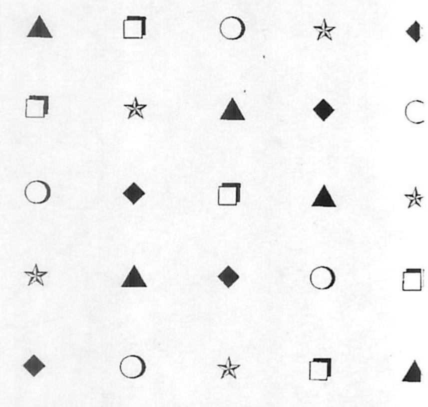

Figure 1. *Each of five symbols appears exactly once in a 5 × 5 latin square.*

Latin squares are easy to come by. Indeed, their number explodes with the size of the square, from two 2×2 squares to twelve 3×3 squares to more than 10^{19} squares of size 8×8. But Dinitz cooked up a variant on the problem of constructing latin squares for which it wasn't clear—until now—that *any* solution could be found.

In an ordinary $n \times n$ latin square, there is only one set of n symbols, and an element from that set must be chosen for each location in the square. In Dinitz's version—called a "partial latin square"—each location is assigned its *own* set of n possible symbols; these sets may vary from location to location. The problem is still to choose a symbol for each location, but now the symbol must come from the set assigned to that location. The goal, however, remains the same: to avoid choosing the same symbol twice in any one row or column.

In Figure 2, a three-element set is assigned to each location in a 3×3 square; the elements in orange constitute a partial latin square. The Dinitz problem asks: Given any assignment of n-element sets of symbols to the n^2 locations in an $n \times n$ array, is it always possible to find a partial latin square? Or to put it negatively, among all the ways to assign n-element sets to the locations of an $n \times n$ array, are there any for which it's impossible to pick an element from each set without picking some symbol twice in the same row or column?

Figure 2. *One symbol (orange) from each three-element set can always be chosen to form a 3 × 3 partial latin square.*

At first glance, the answer seems obvious: Since the problem, in general, uses more than n symbols, it should be easier to satisfy the nonrepetition requirement for a partial latin square than for an ordinary latin square. But that glance overlooks a crucial aspect of the problem: Not every symbol is available at every location. One way to construct an ordinary latin square is to specify where in each row you'll place the first symbol, where the second symbol, and so on; that approach doesn't even make sense for partial latin squares.

Another telling difference between ordinary and partial latin squares casts further doubt on the "obviousness" of the answer. Ordinary latin squares can always be filled in "row by row." If, say, the first two rows of a 5×5 square have been

Jeff Dinitz

filled in successfully (without doubling up in either row or any column), then the rest of the rows can also be filled in to give a latin square. That means that when you're trying to create a latin square, you'll never paint yourself into a corner—you won't get down to the last row, for example, and find yourself unable to complete the square. With partial latin squares, by contrast, you *can* paint yourself in. For example, if the sets in the first row of a 2×2 array are $\{A, B\}$ and $\{B, C\}$, it's natural to choose A and B as the symbols in that row—but then you get in trouble when you see the sets $\{A, C\}$ and $\{B, C\}$ in the next row.

Complications notwithstanding, Dinitz's conjecture—that partial latin squares can always be found—turns out to be true. It just took fifteen years for a proof to be found. In the meantime, the problem served as a kind of drawing card for the theory of combinatorial design and a testing ground for new ideas.

Dinitz's conjecture can be verified directly for 2×2 arrays, because there are so few different possibilities. In principle, the conjecture can be checked for arrays of any given size. That's because there are only finitely many cases to check: The total number of distinct symbols for an $n \times n$ array cannot exceed n^3, so the number of cases is less than n^3 to the power n^3 (more precisely, it's at most the n^2 power of $\binom{n^3}{n}$). But the numbers involved in such a brute-force, case-by-case analysis grow astronomically with n. The 3×3 problem is small enough for this approach to be practical, but the 4×4 case is already out among the stars.

In 1991, however, Noga Alon and Michael Tarsi at Tel Aviv University in Israel proved a theorem that made it easy to verify (by computer) Dinitz's conjecture for 4×4 and 6×6 arrays. Their theorem is not specific to Dinitz's problem. It concerns a general problem in graph theory called "list coloring."

In combinatorics, a graph is a set of points (called *vertices*) and a set of lines or curves (called *edges*) connecting them. Many applications of graphs in scheduling or network theory can be interpreted as coloring the edges of a graph, with the stipulation that no two edges of the same color meet at a common vertex. To schedule a college football season, for example, let each team be represented by a vertex, draw an edge connecting teams that are slated to meet, and then color each edge according to the week on which the two teams are to play (say red for week 1, blue for week 2, and so on). The condition that no like-colored edges should meet at a common vertex simply means that no team should be asked to play two games simultaneously.

In a list-coloring problem, each edge in a graph is assigned a prescribed set, or list, of allowed colors. The Dinitz problem can be viewed as a special case of list coloring, for graphs in which each of n "row" vertices is joined to each of n "column" vertices (see Figure 3). Graphs of this type, in which the vertices are separated into two sets and all edges cross from one set to the other, are known as "bipartite" graphs; the particular graph associated with the Dinitz problem is called a complete bipartite graph, because it includes all possible edges between the two sets of vertices. There is a general conjecture regarding how large the palette of possible colors for each edge of a graph must be in order to ensure that a list coloring is possible. Viewed from this angle, the Dinitz problem is just the tip of an immense theoretical iceberg.

Alon and Tarsi's theorem gives a condition which, if satisfied, guarantees the existence of a list coloring from sets of a particular size. Their condition is simple enough to be verified explicitly for the graphs associated with the 4×4 and 6×6 Dinitz problems. In principle, the condition can be checked for *all* even n, but once again, the amount of computation involved gets quickly out of hand. Furthermore, the condition is *never* satisfied for odd n. (This doesn't mean that the Dinitz conjecture is false for odd n, just that Alon and Tarsi's theorem won't help prove it for those cases.)

Other researchers, notably Roland Häggkvist at the University of Stockholm, had made inroads on the list coloring problem and its relation with the Dinitz conjecture. In late 1992, Jeannette Janssen, then a graduate student at Lehigh University in Bethlehem, Pennsylvania (now a postdoc at Concordia University in Montreal), proved a result that surprised even many of the experts. Janssen showed that Alon and Tarsi's theorem could be used to solve completely a slightly weaker version of Dinitz's problem. Instead of focusing on squares, Janssen looked at *rectangles*—arrays with fewer rows than columns. She showed that in any $r \times n$ array with $r < n$, it's enough to have n symbols (or colors) assigned to each location in order to guarantee that a partial latin rectangle exists.

Janssen's result comes close to the full Dinitz conjecture in two different (but closely related) ways. First, it says that you can always fill in at least the first $n - 1$ rows of a partial latin square (the previous best result guaranteed only two-sevenths of the rows). Second, by starting with an $n \times (n + 1)$ rectangle and then lopping off the last column, Janssen's theorem says that you can always find a

Jeannette Janssen. (Photo courtesy of Cliff Skarstedt.)

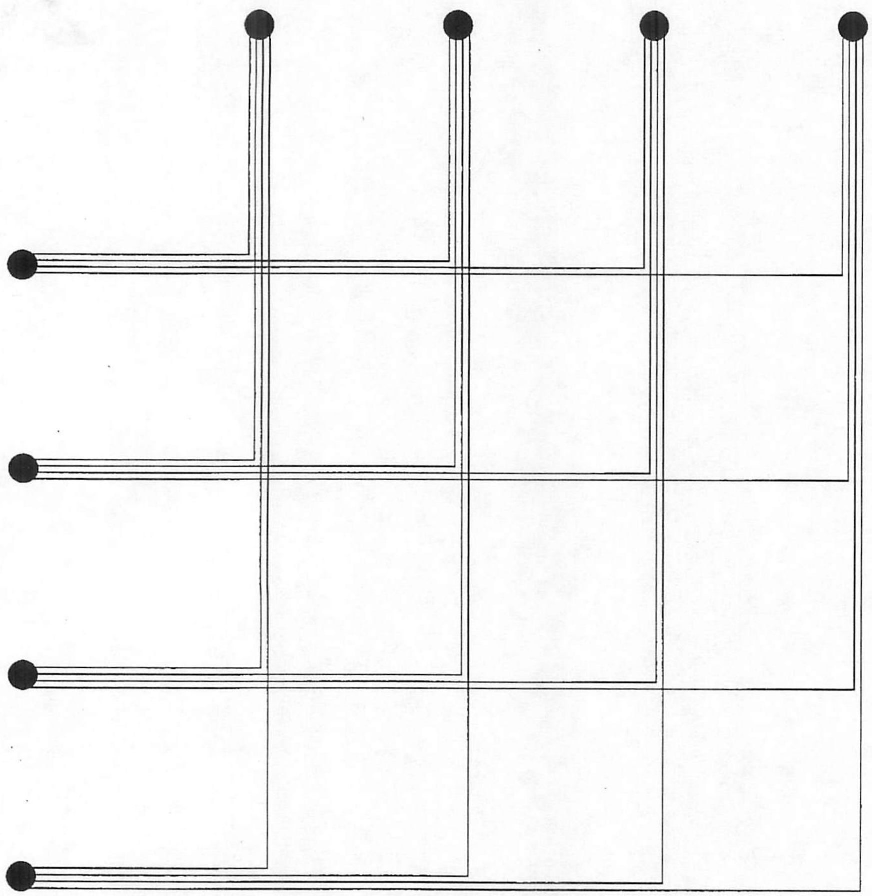

Figure 3. *Each edge in a bipartite graph corresponds to a location in a $n \times n$ array.*

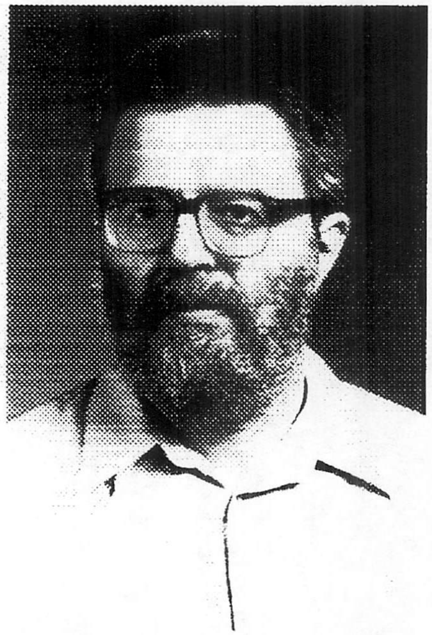

Fred Galvin

partial latin square if $n + 1$ symbols have been assigned to each location—again, far better than previous results.

Experts in the field lauded Janssen's breakthrough. "It is brilliant," said Herb Wilf of the University of Pennsylvania. "It moves the problem much closer to a resolution than anyone had expected." Other theorists agreed, predicting the full Dinitz problem would be solved soon, perhaps within a year. They were right—but not quite for the reasons they had in mind.

Fred Galvin, a mathematician at the University of Kansas, read Janssen's proof in the *Bulletin of the American Mathematical Society*; this led him back to Alon and Tarsi's paper in the journal *Combinatorica*. A remark in that paper made Galvin realize that one of the ideas in Janssen's work could be parlayed into a proof of the complete Dinitz problem, provided one could prove a certain result about the existence of something called a kernel.

Loosely speaking, a kernel of a graph is a "largest possible" subset of vertices, no two of which are connected by an edge. The precise definition is more technical, but the way kernels are used in Galvin's proof is simple: take any color, say red, identify the set of locations that include red among their allowed colors, find a kernel of that set, and then make red your choice for all the locations in that kernel. The Dinitz problem is solved by repeating this process with other colors until every location has been colored—but this approach wouldn't work, Galvin knew, if some set of locations didn't have a kernel.

"I didn't know much about kernels, so I decided to go to the library and see what's available in the way of kernel existence theorems," Galvin recalls. He found exactly what he needed in the second paper he looked at, a theorem by Frédéric Maffray which appeared in the *Journal of Combinatorial Theory (Series B)* in 1992.

"I was really surprised," Galvin says. "I read and reread [Maffray's paper] several times, thinking maybe I misunderstood one of the definitions." That can happen in a technical tangle of terminology—but not this time. Maffray's theorem was indeed the missing ingredient; the Dinitz problem had been solved.

Galvin circulated a three-page, handwritten account of his findings early this year (1994). He subsequently streamlined the proof to make it self-contained. He is still surprised, almost embarrassed, by the proof's simplicity and the way in which he found it. "None of the ideas in the proof originated with me," he says. "All I did was put together a couple of things that were already in the literature."

The experts are also surprised. "The proof is just amazing," says Jeff Kahn, an expert on combinatorics at Rutgers University. Adds Janssen: "Nobody thought that if there would be a proof, it would fit on three pages."

In fact, Galvin's three-page proof solves the list-coloring problem not just for the complete bipartite graphs associated with the Dinitz problem, but for *all* bipartite graphs. Janssen thinks the proof gives insight into the general list-coloring problem for all graphs. Although Galvin's proof uses none of the elaborate theoretical machinery in Alon and Tarsi's paper or in Janssen's work, the heavy-duty stuff may still be crucial in solving the general problem—the Dinitz problem may have turned out easy to solve because it's a special case, Janssen says. On the other hand, the list-coloring problem may ultimately turn out easy to solve as well, perhaps because it's a special case of some even more general problem. If there's a lesson to be drawn, it's that hard problems need not stay that way.