Mathematician Solves Simplex Problem

For decades, investigators have wondered why the simplex algorithm works so well in practice; now there is a proof that it has to

An outstanding problem in economics, operations research, and computer science recently was solved by Stephen Smale, a mathematician at the University of California at Berkeley. The problem—to prove that the enormously popular simplex algorithm must work well for most linear programming problems—has plagued researchers for decades.

Although Smale stresses that his work is still being carefully checked by his colleagues in the mathematics community, those who are familiar with his result are convinced his proof is in essence correct. Richard Cottle of Stanford, for example, says, "On the whole, I would say it looks very good, very exciting." Herbert Scarf of Yale University says, "It's a beautiful piece of work."

The simplex algorithm is one of the most popular computing methods ever devised. First proposed nearly 50 years ago by George Dantzig of Stanford University, it has become essential to businesses world-wide. "It is used by virtually every large company in the U.S.," says B. Curtis Eaves of Stanford.

The reason the simplex algorithm is so popular is that it is the best method available to solve linear programming problems, which are an inescapable part of running large corporations and which also occur in fields such as engineering, biology, agriculture, and the social sciences.

In a typical linear programming problem, a business will want to select its sources of supplies in such a way as to minimize its shipping costs. There will be some constraints on the supplies, such as that all orders be at least a certain size. In its mathematical formulation, the problem is to find the best solution to a set of linear inequalities by considering a function—such as shipping costs—that must be maximized or minimized. Linear programming problems usually involve thousands of inequalities.

These problems, geometrically, are problems of searching the vertices of multidimensional polygons whose boundaries are determined by the constraints on the problems.

The best solution to a linear programming problem always lies on a vertex and the aim of the simplex algorithm is to find that vertex. Since there may be hundreds of thousands of vertices the algorithm would take so much time as to be infeasible, if every vertex had to be checked. But, in practice, the simplex algorithm finds the best solutions to linear programming problems without searching all vertices. As Dantzig explains, thousands or even tens of thousands of linear programming problems are routinely solved each day with the simplex algorithm. However, computer scientists and mathematicians can easily devise problems that the method simply cannot solve because, in attempting to solve them, the algorithm would check every vertex.

"There is a disparity," says Eaves. "In 30 years of running the simplex algorithm, it has given outstanding results. But, theoretically, there are problems that can make the simplex look as bad as any algorithm." So an outstanding question has been, is it possible to prove that, on the average, the simplex algorithm must work well?

Smale's work should be a great stimulus to other researchers, says Eaves.

"Average" behavior can be notoriously difficult to quantify. Often it is easier to do a worst-case analysis because it is so hard to decide what makes an average problem average. Yet, says Cottle, worst-case analysis is not of prime interest. "It is the efficiency of the simplex in the real world that people are interested in understanding," he remarks.

What Smale shows is that there is an upper bound on the number of computational steps necessary to solve the average linear programming problem. The upper bound is proportional to the number of variables in the problem if the number of constraints is fixed. "It's a remarkable bound. It's better than I would have expected," Scarf remarks.

The key to proving that the simplex algorithm works well on the average, Smale says, was to "put a probability measure on the space of data" to define average behavior and then to find the right geometric way to look at the problem. He drew on a body of work, dating back to the 1960's and attributable chiefly to Dantzig, Cottle, and Carl Lemke of Rennselaer Polytechnic Institute. Smale also made use of more recent material and, Cottle points out, "He made a few clever observations that made all the difference in the world."

"I look at linear programming in terms of certain path-following methods," Smale explains. In two dimensions, the problem is like looking at a special set of cones in the plane with a straight line, drawn at random, passing through some of them. Each time the line moves across a cone, one step is taken in the simplex algorithm. The problem is to decide how frequently the line encounters new cones. With Smale's analysis, this problem is one of determining the average area of the cones rather than determining how they are arranged. In n dimensions, it is a problem of determining the average volume of n-dimensional cones.

In his analysis of the problem, Smale relies on an analog of the global Newton method which he and others first began developing about 6 years ago. This method is similar to piecewise linear homotopy methods (*Science*, 4 May 1979, p. 488) which are often used for finding fixed points of economic equilibrium. In these homotopy methods, mathematicians take a problem whose answer they want and a problem whose answer they know. They connect the two with a family of problems and thereby are able to estimate the desired answer.

In the global Newton methods, says Smale, "I deal with just one problem, such as $f(x) = y$. I know y and I want to find x, so I take two points, x_0 and y_0. I know $f(x_0) = y_0$. Then I take a path from y_0 to y and lift it back to a path that goes from x_0 to x. That gives me a solution to the original equation."

Eventually, Smale predicts, his sort of analysis may show why the simplex method does so badly on problems designed to foil it. As for now, says Eaves, Smale's work should be a great stimulus to other researchers. "With little doubt, Steve's paper will generate 50 or so papers in the next year to follow up," Eaves remarks.—GINA KOLATA

0036-8075/82/0702-0039$01.00/0