



THE (DANTZIG) SIMPLEX METHOD FOR LINEAR PROGRAMMING

George Dantzig created a simplex algorithm to solve linear programs for planning and decision-making in large-scale enterprises. The algorithm's success led to a vast array of specializations and generalizations that have dominated practical operations research for half a century.

George Dantzig, in describing the “Origins of the Simplex Method,”¹ noted that it was the availability of early digital computers that supported and invited the development of LP models to solve real-world problems. He agreed that invention is sometimes the mother of necessity. Moreover, he commented that he initially rejected the simplex method because it seemed intuitively more attractive to pursue the objective function downhill—as in the currently popular interior-point methods—rather than search along the constraint set’s edges. It is this latter approach that the simplex method uses, which should not be confused with the J.A. Nelder and R. Mead’s function-minimization method,² also associated with the word simplex.

When Dantzig introduced his method in 1947, it was somewhat easier to sort out the details of the simplex method than to deal with the “where are we in the domain space” questions that are, in my opinion, the core of interior-point approaches. Easier does not mean simpler, however.

Chapter and verse

The LP problem is, in one of the simplex method’s many forms,

$$\begin{aligned} \text{Minimize (with respect to } \mathbf{x}) \quad & \mathbf{c}^T \mathbf{x} && (1a) \\ \text{subject to} \quad & \mathbf{A} \mathbf{x} = \mathbf{b} && (1b) \\ \text{and} \quad & \mathbf{x} \geq 0 && (1c) \end{aligned}$$

This is not always the problem we want to solve, though—we might want to have the matrix \mathbf{A} and vector \mathbf{b} partitioned row-wise so that

$$\mathbf{A} = \begin{matrix} (\mathbf{A1}) \\ (\mathbf{A2}) \\ (\mathbf{A3}) \end{matrix} \quad \mathbf{b} = \begin{matrix} (\mathbf{b1}) \\ (\mathbf{b2}) \\ (\mathbf{b3}) \end{matrix}$$

so that we can write an LP problem as

$$\text{Minimize (with respect to } \mathbf{x} \text{)} \quad \mathbf{c}'\mathbf{x} \quad (2a)$$

$$\text{subject to} \quad \mathbf{A1} \mathbf{x} \leq \mathbf{b1} \quad (2b)$$

$$\mathbf{A2} \mathbf{x} = \mathbf{b2} \quad (2c)$$

$$\mathbf{A3} \mathbf{x} \geq \mathbf{b3} \quad (2d)$$

$$\mathbf{x} \geq 0 \quad (2e)$$

However, the use of slack and surplus variables (the \mathbf{x} variables we use to augment our problem) can transform the problems from one to the other. In fact, we must add slacks to the left-hand side of Equation 2b and subtract them from Equation 2d to arrive at equalities. Moreover, the special non-negativity conditions in Equations 1c and 2e could be subsumed into the matrix equation or inequation structure, but the LP tradition calls for listing them separately. The number of rows in \mathbf{A} apart from the non-negativity constraints we call m , and the number of variables (including the slacks and surpluses) we call n . Typically, $m \leq n$.

The simplex method assumes we have an initial basic feasible solution \mathbf{x}_{init} , that is, a solution that is feasible and that has just m of the \mathbf{x} 's non-zero. The $m\mathbf{x}$ values multiply m columns of \mathbf{A} , which we call the basis. "Basic" definitely has a technical rather than a general meaning here. We assume this set of basis vectors (columns of \mathbf{A}) is linearly independent and that it has full rank. Therefore, the core of the simplex method is to exchange one of the columns of \mathbf{A} that is in the basis for one that is not. This corresponds to increasing one of the nonbasic variables while keeping the constraints satisfied, a process that reduces some of the \mathbf{x} 's. We choose to increase the \mathbf{x} that gives us the most decrease in the objective function, and we continue to increase it until one of the current basic variables goes to zero, which means we now have a new basis. Moreover, the new solution turns out to be a new vertex of the simplex that the constraints define, and it is also a neighboring vertex to that described by our starting basic feasible solution. We have exchanged one vertex of the simplex for one of its neighbors and done so in a way that moves the objective function toward the minimum.

This exchange, or pivoting, seems to be a topic especially devised to torture business school undergraduates. Nevertheless, it is theoretically and practically a relatively simple process, unless, of course, you have a computer without enough memory to hold everything all at once. Additionally, there are some nasty details we have to sort out:

- How do we get an initial feasible solution? Or, more importantly, how do we tell if there is no feasible solution?
- What happens if the constraints don't stop the objective from being reduced? That is, what if the solution is unbounded?
- What happens if several vertices give equal objective values?

These questions, and the memory-management issues, occupied many researchers and generated many journal and book pages over the years. For example, more than half the pages in S. Gass's classic text are devoted to such details;³ the applications come so late in the book that students of slow lecturers must have wondered what all the fuss was about.

Researchers handled the questions mentioned earlier in a variety of ways. Clever management and programming tactics overcame the memory issue. In early years, storage was so slow that programmers spent considerable effort to arrange computations that would complete in time to read and write data as disks and drums presented the appropriate tracks and sectors. Such complications help to obscure the central ideas of the algorithms in a muddle of detail, mostly irrelevant now.

The initial feasible solution issue was addressed, perhaps surprisingly, by adding m more variables, called artificial variables, which form the initial basis, then minimizing an objective function that drives them all out of the basis.⁴ If we cannot drive them out, we have an infeasible problem.

Unbounded solutions turn out to be rather simple to detect, but the "equal objective function" or degeneracy issue worried people a great deal, because they implied the possibility that the algorithm might not terminate. In practice, perturbation or rule-based methods can avoid the cycling.

Nice extras

Modern textbooks spend much less ink on the details.⁵ Similarly, recent discussions⁶ of the computational techniques have abandoned the tableaux and the jargon of earlier work^{3,4} regarding the matrix notations of numerical linear algebra. What remains of great interest, both mathematically and computationally, is the duality feature of LP problems. That is, an LP written as

$$\text{(Primal) minimize } \mathbf{c}'\mathbf{x} \text{ such that } \mathbf{A}\mathbf{x} \geq \mathbf{b}, \mathbf{x} \geq 0$$

has a dual equivalent

(Dual) maximize $\mathbf{b}'\mathbf{y}$ such that $\mathbf{A}'\mathbf{y} \leq \mathbf{c}$, $\mathbf{y} \geq 0$.

Moreover, the optima, if they exist, have the same objective value. Duality also leads to useful interpretations of solutions in terms of prices or values when the LP problems under consideration have economic or related contexts.

Practical importance

The simplex method's importance really lies in the value of the LP applications, even when the LP model is only a crude approximation to the real world. In the 1940s, many organizations were very hungry for solutions to LP problems, even if they did not realize what LP was. Oil and chemical companies led the way, especially for optimizing product mix from multiple sources or multiple sites. Transportation companies and the military recognized quite early that transportation and logistics problems could be formulated as LPs. Large-scale agricultural economic problems were also early applications. Examples also became realized in production planning, staff and resource scheduling, and network or traffic flows. LP can even be applied to maintaining the confidentiality of government statistics.⁷

For example, the Diet Problem has as objective function elements (\mathbf{c}) the costs per unit of food ingredients. The right-hand side values (\mathbf{b}) are the required minimum amounts of each of a list of m nutrients, and the constraint coefficients \mathbf{A} give the amount of each nutrient in a unit of each ingredient. Thus we want to have $\mathbf{A}\mathbf{x} \geq \mathbf{b}$ to satisfy the nutrient requirements (we could also put in upper limits of nutrients like vitamin A that can be toxic), but we also want the cheapest blend. Of course, the resulting solution might not be very tasty. It is simply nutritious, and by construction will be the cheapest such recipe. Perhaps this is the origin of institutional cafeteria food.

Although the mathematics and computational algorithms might be fascinating, my view is that the value of the applications is the first reason for the importance of the simplex method. Promotion to the "top 10" category still needs something

else, and this is the particular efficiency of the simplex method in finding the best vertex of the simplex. Most constraint matrices in practical LP problems are quite sparse, and the number of iterations where we move from one vertex to another is generally very small relative to the number of variables \mathbf{x} . Even better, we can explore the sensitivity of the optimal solution to small changes in the constraint and objective very easily. This is not always the case for some other methods.

Did Dantzig realize this efficiency as he first coded the simplex method? My guess is that he did not, although it is clear from his reminiscences¹ that such an understanding was not long in coming. In looking back over half a century, I find it remarkable that so many workers could sort through the jungle of awkward, and to some extent unnecessary, details to see the underlying tool's value. \square

References

1. S.G. Nash, *A History of Scientific Computing*, ACM Press, New York, 1990, pp. 141–151.
2. J.A. Nelder and R. Mead, "A Simplex Method for Function Minimization," *Computer J.*, Vol. 7, 1965, pp. 308–313.
3. S. Gass, *Linear Programming*, 2nd ed., McGraw-Hill, New York, 1964.
4. G. Hadley, *Linear Programming*, Addison-Wesley, Reading, Mass., 1962.
5. S.G. Nash and A. Sofer, *Linear and Nonlinear Programming*, McGraw-Hill, New York, 1996.
6. J.L. Nazareth, *Computer Solution of Linear Programs*, Oxford Science Publications, New York, 1987.
7. G. Sande, "Automated Cell Suppression to Preserve Confidentiality of Business Statistics," *Statistical J. United Nations ECE*, Vol. 2, 1984, pp. 33–41.

John C. Nash is a professor at the University of Ottawa. He is the author or coauthor of three books on computation, and his research and writings cover a wide range of topics in computers, mathematics, forecasting, risk management, and information science. He has been a mathematics columnist for *Interface Age* and the scientific computing editor for *Byte* magazine. He obtained his BSc in chemistry at the University of Calgary and his DPhil in mathematics from Oxford. Contact him at the Faculty of Administration, Univ. of Ottawa, 136 J-J Lussier Private, Ottawa, Ontario, K1N 6N5, Canada; jcnash@uottawa.ca; macnash.admin.uottawa.ca.