

The Simplex Method of Linear Programming Using LU Decomposition

RICHARD H. BARTELS AND GENE H. GOLUB Stanford University,* Stanford, California

Standard computer implementations of Dantzig's simplex method for linear programming are based upon forming the inverse of the basic matrix and updating the inverse after every step of the method. These implementations have bad round-off error properties. This paper gives the theoretical background for an implementation which is based upon the LU decomposition, computed with row interchanges, of the basic matrix. The implementation is slow, but has good round-off error behavior. The implementation appears as CACM Algorithm 350.

KEY WORDS AND PHRASES: simplex method, linear programming, LU decomposition, round-off errors, computational stability CR CATEGORIES: 5.41

1. LU Decomposition

The linear programming problem:

maximize
$$d^{T}x$$

subject to $\begin{cases} Gx = b \\ x \ge 0 \end{cases}$ (1)

where $d^{T} = [d_{0}, \dots, d_{n-1}], b^{T} = [b_{0}, \dots, b_{m-1}]$, and

$$G = \begin{bmatrix} g_{0,0} & \cdots & g_{0,n-1} \\ \vdots & & \vdots \\ g_{m-1,0} & \cdots & g_{m-1,n-1} \end{bmatrix}$$

are given, with $b \ge 0$, is commonly solved by a two-phase process built upon the simplex method of G. B. Dantzig [2]. Each step of this method requires the solution of three systems of linear equations involving a common matrix of coefficients, the basis matrix P. It is the usual practice to solve these systems by forming P^{-1} , either directly or in

* Computer Science Department. This project was supported in part by NSF under project GP948 and ONR under project NR 044 211. factored form as a product of transformations, and then applying it to the right-hand sides of the systems.

The basis matrix of any simplex step differs from that of the preceding step in only one column, so it is easier to update P^{-1} than to invert the new P. While this generally produces satisfactory results, it is vulnerable to computational problems in two respects. First, if P^{-1} is continually updated, computational inaccuracies imposed by limitedprecision arithmetic (round-off errors) are allowed to propagate from step to step. Second, the updating, in whatever way it is carried out, is equivalent to premultiplying P^{-1} by a matrix of the form:

	1	0	$-y_0/y_k$		
	•		•	0	
	0	1	$- y_{k-1}/y_k$		
k	0		$1/y_k$	0	
			$- y_{k+1}/y_k$.	1.	0
	0				
			$- y_{m-1}/y_k$	0	1
			k		

In the presence of round-off errors this computation can be quite inaccurate if y_k is small relative to the other y_j .

Both of these problems can be eliminated if, instead of P^{-1} , the LU decomposition of P is computed using row interchanges to select pivots. The problem of solving a linear system based upon P is then reduced to that of backsolving two triangular linear systems. The numerical stability of this scheme is generally quite good (see [4]). Advantage can be taken of the similarity between any two successive matrices so as to economize on the computation of the LU decompositions. For details see [1, Secs. 5–6].

Additional accuracy is obtained in the program (see Algorithm 350) by iteratively refining the solution to problem (1), after it has been found, according to the scheme given in [4, p. 121].

The algorithm requires $n^3/6 + O(n^2)$ multiplications per exchange on the average. Thus the algorithm is most effectively used for very ill-conditioned problems or in conjunction with other implementations of the simplex algorithm which produces an initial basis more cheaply. A fast and accurate version of the simplex algorithm is described in [5] but this requires additional storage.

This paper gives the theoretical background of Algorithm 350, "Simplex-Method Procedure Employing LU Decomposition," by the same authors, which appears on pages 275–278.

2. The Two-Phase Solution Process

The possibility of degeneracy is ignored, as is customary. PHASE 1 (Find a basic feasible solution.)

maximize
$$-e^{T}a$$

subject to
$$\begin{cases} Gx + a = b \\ x \ge 0, \quad a \ge 0, \end{cases}$$
 (2)

where $a^{T} = [a_{0}, \dots, a_{m-1}]$ is a vector of "artificial variables" and $e^{T} = [1, \dots, 1]$ (*m* components). The solution is obtained by starting with a = b as a basic feasible solution and applying the simplex method. As each a_{i} becomes nonbasic, it may be dropped from the problem altogether. Since $-e^{T}a \leq 0$ for all a, problem (2) has a solution. If $max(-e^{T}a) \neq 0$, then problem (1) has no basic feasible solution. Otherwise, the second phase is entered.

PHASE 2 (Find an optimal basic feasible solution.)

If the solution to problem (2) contains only components of x as basic variables, these become the initial basic variables for solving problem (1) by the simplex method. If artificial variables, say a_{j_0} , \cdots , a_{j_l} , remain basic but at zero level in the solution of (2), an equivalent problem to (1) is solved:

maximize $d^T x$



where e_k represents the kth column of the identity matrix (column numbering begun at zero). The additional variable s, together with the additional constraint

$$s+\sum_{i=0}^{l}a_{j_i}=0,$$

holds the artificial variables at zero level throughout further computation. The simplex method is applied to problem (3) with s and the basic variables from problem (2) as the initial set of basic variables. (For an example of this two-phase process in operation see [3, Sec. 7.7].)

3. Some Computational Details

Phase 1 is skipped if a basic feasible set of variables for problem (1) is known a priori. If a partial set of basic variables is known for problem (1), the computation of phase 1 can be restructured to take this into account. To be precise, let κ ($1 \le \kappa \le m - 1$) columns of G (without loss of generality these can be made $g_{n-\kappa}, \dots, g_{n-1}$, the last

columns of G) be used to construct the following matrix:

P =	e ₀		$e_{m-\kappa-1}$	gn-к		g_{n-1}	} .
-----	----------------	--	------------------	------	--	-----------	-----

Suppose that $P^{-1}b \ge 0$. (This situation arises if the variables are numbered so that $x_{n-\kappa}$, \cdots , x_{n-1} are the slack variables and the constraints are ordered so that $g_{n-\kappa} = e_{m-\kappa}$, \cdots , $g_{n-1} = e_{m-1}$, whereby P becomes the identity matrix.) Then problem (2) can be replaced by the problem:

maximize $-\tilde{e}^T \tilde{a}$

subject to
$$\begin{cases} \boxed{\begin{array}{c|c} G^{(1)} & \tilde{I} \\ \hline & 0 \\ \end{array}} & G^{(2)} \\ \hline & & \\ \hline & & \\ \hline & & \\ x^{(2)} \\ \end{array} \\ x^{(1)} \ge 0, \quad x^{(2)} \ge 0, \quad \tilde{a} \ge 0 \end{cases}$$
(4)

where $x^{(1)} = [x_0, \dots, x_{n-\kappa-1}], x^{(2)} = [x_{n-\kappa}, \dots, x_{n-1}],$ $G^{(1)}$ and $G^{(2)}$ are appropriate submatrices of G, $\tilde{a} = [a_0, \dots, a_{m-\kappa-1}], \tilde{e} = [1, \dots, 1] (m - \kappa \text{ components}), \text{ and } \tilde{I}$ is the $(m - \kappa)$ -order identity matrix. The simplex method is begun with $a_0, \dots, a_{m-\kappa-1}, x_{n-\kappa}, \dots, x_{n-1}$ as basic variables. Since there are fewer artificial variables to be driven to zero in problem (4) than in problem (2), we expect problem (4) to require on the average less work to solve.

A reduction of computation beyond that described in [1] can be arranged for the LU decompositions calculated in phase 1. Each time an artificial variable becomes nonbasic, a row and column interchange can be applied to the basis matrix so that it always has the partitioned form

$$P = \boxed{\begin{array}{c|c} I \\ \hline O \\ \hline \end{array}} \boxed{\begin{array}{c} R \\ \hline S \\ \hline \end{array}}$$

where I is the identity matrix of appropriate order, and R and S are composed of components of G. The LU decomposition of P is known when the decomposition $S = \mathcal{L}$ has been computed; viz.

D	Ι	0	I	R
1 =	0	L	0	น

The triangular systems of equations induced by this format are solved in a particularly efficient manner by Algorithm 350.

4. Example Application

Let V be a random variable restricted to a finite set of values v_0 , \cdots , v_{n-1} . Let p_i be the probability that V has value v_i . If certain moments

$$\mu_{k} = \sum_{j=0}^{n-1} v_{j}^{k} p_{j}, \qquad k \in \{k_{0}, \cdots, k_{m-1}\}$$

are known together with the values v_j , estimates of the

 p_i and of unknown moments μ_l can be made via linear programming; viz.

where

$$\alpha_i \leq p_i \leq \beta_i \quad ext{and} \quad \gamma_l \leq \mu_l \leq \delta_l$$
,

$$\alpha_i = -\max(-x_i), \qquad \beta_i = \max(x_i)$$

$$\gamma_l = -\max(-\sum_{j=0}^{n-1} v_j^l x_j), \qquad \delta_l = \max(\sum_{j=0}^{n-1} v_j^l x_j),$$

all subject to the constraints

$$egin{aligned} &x_0 \geq 0, \ \cdots, \ x_{n-1} \geq 0, & \sum_{j=0}^{n-1} x_j = 1, \ &\sum_{j=0}^{n-1} v_j^k x_j = \mu_k \ , & k \in \{k_0, \ \cdots, \ k_{m-1}\}, \ k
eq 0. \end{aligned}$$

The basis matrices encountered in solving these problems are submatrices of a Vandermonde matrix, making them somewhat poorly conditioned; so it is important that the simplex method be carried out accurately. Only the objective function differs from one estimation problem to the next; so the solution of any one of the problems would provide an initial basis for the other problems. Phase 1 would need to be carried out only once.

Thus, for example, to compute upper bounds for p_0, \dots, p_{n-1} from given values v_0, \dots, v_{n-1} and given moments μ_0, \dots, μ_{m-1} , one would set up arrays $G[i, j] \equiv v_i^i$ and $b[i] \equiv \mu_i$ and execute

 $\begin{aligned} kappa &:= 0; \quad \text{for } i := 0 \text{ step 1 until } n-1 \text{ do } d[i] := 0; \\ \text{for } i &:= 0 \text{ step 1 until } n-1 \text{ do} \\ \text{begin } d[i] &:= 1; \\ & \text{if } i \geq 1 \text{ then } d[i-1] := 0; \\ linprog(m, n, kappa, G, b, d, x, z, ind, infeas, unbdd, sing); \\ upper bound[i] &:= z; \\ \text{end}; \end{aligned}$

As an illustration, the preceding was run on Stanford University's B5500 computer with the data $\mu_0 = 1$, $\mu_1 = 3$, $\mu_2 = 10.5$, $\mu_3 = 40.5$ and $v_i = i$ for $i = 0, \dots, 6$. The results are listed below.

Computed upper bound	Actual probabil- ily
.083333333325	.015625
.19999999998	.09375
.5000000001	.234375
.83333333335	.3125
.499999999999	.234375
.20000000001	.09375
.083333333334	.015625

Lower bounds for the p_i computed from the above data in a similar fashion were all zero. Computed bounds for μ_4 through μ_{10} were as follows.

Computed lower bound	Actual momeni	Computed upper bound
165.5	168.0	175.5
700.5	738.0	850.5
3035.5	3396.75	4495.5
13390.5	16251.75	25150.5
59965.5	80335.5	145435.5
272200.5	408280.5	856210.5
1251135.5	2124764.25	5088055.5

RECEIVED SEPTEMBER, 1967; REVISED NOVEMBER, 1968

REFERENCES

- 1. BARTELS, R. H., AND GOLUB, G. H. Stable numerical methods for obtaining the Chebyshev solution to an overdetermined system of equations. Comm. ACM 11, 6 (June 1968), 401-406.
- DANTZIG, G. B. Linear Programming and Extensions. Princeton U. Press, Princeton, N.J., 1963.
- 3. HADLEY, G. Linear Programming. Addison-Wesley, Reading, Mass., 1962.
- 4. WILKINSON, J. H. Rounding Errors in Algebraic Processes. Prentice-Hall Series in Automatic Computation, Englewood Cliffs, N.J., 1963.
- 5. BARTELS, RICHARD H. A numerical investigation of the simplex method. Tech. Rep. No. CS 104, Comput. Sci. Dep., Stanford U., Stanford, Calif., July 1968.

Rough and Ready Error Estimates in Gaussian Integration of Analytic Functions

PHILIP RABINOWITZ

The Weizmann Institute of Science,* Rehovoth, Israel

KEY WORDS AND PHRASES: numerical integration, analytic functions, error estimates, Gaussian integration, tabulated error coefficients, computable error coefficients, Cauchy integral formula, Chebyshev polynomials CR CATEGORIES: 5.11, 5.16

* Department of Applied Mathematics

268 Communications of the ACM

Two expressions are derived for use in estimating the error in the numerical integration of analytic functions in terms of the maximum absolute value of the function in an appropriate region of regularity. These expressions are then specialized to the case of Gaussian integration rules, and the resulting error estimates are compared with those obtained by the use of tables of error coefficients.

Davis and Rabinowitz [1] presented a method for estimating the error committed in integrating an analytic function by an arbitrary integration rule

$$I(f) = \sum_{i=1}^{n} w_i f(x_i).$$

If the error is denoted by $E(f) = \int_{-1}^{1} f(x) dx - I(f)$, we

Volume 12 / Number 5 / May, 1969