

Notes on the Dual Simplex Method

0. The standard dual simplex method
1. A more general and practical dual simplex method
2. Phase I for the dual simplex method
3. Degeneracy in the dual simplex method
4. A generalized ratio test for the dual simplex method

0. The standard dual simplex method can be derived by considering a linear program of the form

$$\begin{aligned} (\mathbf{D}_0) \quad & \text{Maximize} \quad \pi b \\ & \text{Subject to} \quad \pi A \leq c \end{aligned}$$

where A is an $m \times n$ coefficient matrix. This LP has a natural geometric interpretation, wherein the constraints of $\pi A \leq c$ define half-spaces that intersect in a polytope, and the maximum of πb must occur (if finite) at one or more vertices of the polytope. Starting at some vertex, the dual simplex method moves from vertex to vertex along edges of the polytope, in such a way that each vertex has a higher value of πb than the preceding one.¹ When no such move can be made, the method stops and declares the current vertex optimal.

This geometric view is readily converted to an algebraic statement of the necessary computations. A vertex is represented by a basic feasible solution, which is defined by the intersection of some m independent “tight” constraints. Specifically, letting B be an $m \times m$ matrix of linearly independent columns from A , and c_B be the corresponding subvector of c , a basic solution $\bar{\pi}$ is defined by $\bar{\pi}B = c_B$, and is feasible if in fact $\bar{\pi}A \leq c$. We write \mathcal{B} for the set of basic indices from $\{1, \dots, m\}$, and \mathcal{N} for complementary set of nonbasic indices.

Given a basic feasible solution, an edge radiating from the associated vertex is defined by relaxing some tight constraint $i \in \mathcal{B}$, while all of the other tight constraints remain tight. In algebraic terms, we have a ray defined by $\pi(\theta_i)B = c_B + \theta_i e_i$ for $\theta_i \leq 0$,² and the edge contained in this ray can be expressed as

$$\pi(\theta_i) = \bar{\pi} + \theta_i e_i B^{-1}, \quad 0 \geq \theta_i \geq \Theta_i.$$

Along this edge, the objective value is $\pi(\theta_i)b = \bar{\pi}b + \theta_i e_i B^{-1}b$. If we define $\bar{x}_B = B^{-1}b$ to be a vector whose components are \bar{x}_i , $i \in \mathcal{B}$, then $\pi(\theta_i)b = \bar{\pi}b + \theta_i \bar{x}_i$, and it is easy to see that the objective increases along the edge if and only if

$$\bar{x}_i < 0, \quad \text{where } B\bar{x}_B = b.$$

The dual simplex method chooses to relax some tight constraint $q \in \mathcal{B}$ such that $\bar{x}_q < 0$, or else declares optimality if $\bar{x}_B \geq 0$.

It remains to determine what new vertex is reached as we move along the edge. Writing a_j for the j th column of A , and substituting the expression for $\pi(\theta_q)$ above, we can see that a non-tight constraint $j \in \mathcal{N}$ becomes tight when

$$\pi(\theta_q)a_j - c_j = (\bar{\pi}a_j - c_j) - \theta_q(\sigma_q a_j) = 0, \quad \text{where } \sigma_q B = -e_q.$$

The vertex is reached when the first such constraint becomes tight; that is, when θ_q equals

$$\Theta_q = \max_{\substack{j \in \mathcal{N} \\ \sigma_q a_j > 0}} \frac{\bar{\pi}a_j - c_j}{\sigma_q a_j}.$$

¹For now we ignore the possibility of degenerate iterations, to keep the development simple. Later it will be seen that the dual simplex method may handle certain degenerate situations in an especially effective way.

²Ordinarily we would arrange to have $\theta_i \geq 0$ in this derivation, but it will be seen that $\theta_i \leq 0$ is more natural for generalization to other dual formulations.

The dual simplex method chooses some index p at which this minimum is achieved, and constraint p joins the set of tight constraints. That is, p joins \mathcal{B} while q joins \mathcal{N} .

To finish the computational work of the iteration, some factorization of the matrix B is updated to reflect the replacement of a_q by a_p . The “slack” values $\bar{s}_j = \bar{\pi}a_j - c_j$ are updated by setting \bar{s}_q to Θ_q , and

$$\bar{s}_j \leftarrow \bar{s}_j - \Theta_q(\sigma_q a_j), \quad j \in \mathcal{N} \setminus \{q\}.$$

The vector $\bar{\pi}$, although used in the derivation, is not needed for any step of the method. Thus the principal work of an iteration involves solving one system in B (for \bar{x}),³ solving one system in B^T (for σ_q), and computing the inner products $\sigma_q a_j$ for all $j \in \mathcal{N}$.

Of course the matrix B is the same as the basis matrix for the associated primal linear program,

$$\begin{aligned} (\mathbf{P}_0) \quad & \text{Minimize} && cx \\ & \text{Subject to} && Ax = b \\ & && x \geq 0 \end{aligned}$$

An iteration of the standard primal simplex method also solves one system in B and one in B^T , and forms inner products with the columns of A . As a result, the same data structures and linear algebra routines are used for both the primal and dual methods in simplex implementations.

In practice, the primal LP has arbitrary bounds on the variables, and the dual method must be able to begin from an infeasible basic solution. We will show that the preceding derivation, working only in the geometrically intuitive space of the π variables, can be extended to produce a practical dual simplex method.

1. We consider a particular generalization of the dual linear program that may be written concisely as

$$(\mathbf{D}) \quad \text{Maximize} \quad \pi b - (\pi A - c)u^\pm$$

The objective here is to be interpreted as

$$\sum_{i=1}^m \pi_i b_i - \sum_{j=1}^n (\pi a_j - c_j) u_j^\pm,$$

where

$$u_j^\pm = \begin{cases} u_j^- & \text{if } \pi a_j - c_j < 0 \\ u_j^+ & \text{if } \pi a_j - c_j > 0 \end{cases}, \quad u_j^- \leq u_j^+.$$

This is a problem of maximizing a concave piecewise-linear function that has a certain special form. We’ll first present a dual simplex method for this problem— one that differs only in details from the standard method described above. Then we’ll show that this version of the method exhibits all of the features to be found in large-scale dual simplex implementations.

³The vector \bar{x} may be updated rather than recomputed at each iteration, but still a linear system in B must be solved.

In analyzing this generalized formulation, we can still think of the intersections of m linearly independent hyperplanes $\pi a_j = c_j$ as “vertices”, and can still regard the line segments between vertices as “edges”. In this case, however, every intersection of m independent hyperplanes is feasible; the vertices and edges are not extreme points and edges of the feasible region, but of various convex regions of linearity. Nevertheless it is still true that, if (\mathbf{D}) has a finite optimal value, then it achieves that value at some vertex. Moreover, an optimal vertex can be found by local improvement, moving from vertex to vertex along edges.⁴

Since vertices for (\mathbf{D}) are defined the same as before, so are basic feasible solutions. Letting B be an $m \times m$ matrix of linearly independent columns from A , and c_B be the corresponding subvector of c , a basic feasible solution $\bar{\pi}$ is defined by $\bar{\pi}B = c_B$; no basic solutions are infeasible in this case. Given such a basic feasible solution, we also define for each nonbasic index $j \in \mathcal{N}$ the value

$$\bar{u}_j = \left\{ \begin{array}{ll} u_j^- & \text{if } \bar{\pi}a_j - c_j < 0 \\ u_j^+ & \text{if } \bar{\pi}a_j - c_j > 0 \end{array} \right\},$$

which is simply the slope of the objective term $(\pi a_j - c_j)u_j^\pm$ in the neighborhood of $\bar{\pi}$.⁵

Given a basic solution, an edge radiating from the associated vertex is defined by relaxing the basic equation corresponding to some index $i \in \mathcal{B}$, while all of the other basic equations remain tight. But in this case, since all edges are feasible, the equation may be relaxed in either of two directions. In algebraic terms, we have a line defined by $\bar{\pi}B = c_B + \theta_i e_i$, and the edges contained in this line can be expressed as

$$\begin{aligned} \pi(\theta_i) &= \bar{\pi} + \theta_i(e_i B^{-1}), \quad 0 \geq \theta_i \geq \Theta_i^- && \text{(the negative edge),} \\ \pi(\theta_i) &= \bar{\pi} + \theta_i(e_i B^{-1}), \quad 0 \leq \theta_i \leq \Theta_i^+ && \text{(the positive edge).} \end{aligned}$$

Substituting $\pi(\theta_i)$ for π in the objective function of (\mathbf{D}) , we see that the change as a function of θ_i now has three terms:

$$\begin{aligned} \theta_i e_i B^{-1} b &&& \text{from } \pi b \text{ (as in §0)} \\ -\theta_i u_i^\pm &&& \text{from } -(\pi a_i - c_i)u_i^\pm \\ -\sum_{j \in \mathcal{N}} \theta_i e_i B^{-1} a_j \bar{u}_j &&& \text{from } -(\pi a_j - c_j)u_j^\pm, \text{ for all } j \in \mathcal{N} \end{aligned}$$

Adding these together, we see that the overall change in the objective is

$$\theta_i \left[-u_i^\pm + e_i B^{-1} (b - \sum_{j \in \mathcal{N}} a_j \bar{u}_j) \right].$$

The derivation can now proceed much as it did in §0. If we define $\bar{x}_B = B^{-1}(b - \sum_{j \in \mathcal{N}} a_j \bar{u}_j)$ to be a vector whose components are \bar{x}_i , $i \in \mathcal{B}$, then the change in the

⁴These properties of (\mathbf{D}) can be derived as corollaries of well-known properties of convex separable piecewise-linear minimization. Alternatively, (\mathbf{D}) may be converted to a linear program of the form (\mathbf{D}_0) by adding one variable, such that each basic point in (\mathbf{D}) has a corresponding extreme point in (\mathbf{D}_0) ; then the result can be derived from the usual properties of basic solutions for LPs.

⁵If $\bar{\pi}a_j = c_j$ then \bar{u}_j can be taken as either u_j^- or u_j^+ . We'll return to this degenerate case in §3.

objective is simply $\theta_i[-u_i^\pm + \bar{x}_i]$, and the objective is seen to increase

$$\left. \begin{array}{l} \text{along the negative edge, if } \bar{x}_i < u_i^- \\ \text{along the positive edge, if } \bar{x}_i > u_i^+ \end{array} \right\} \text{ where } B\bar{x}_B = b - \sum_{j \in \mathcal{N}} a_j \bar{u}_j.$$

The dual simplex method thus chooses either to relax negatively some equation $q \in \mathcal{B}$ such that $\bar{x}_q < u_q^-$, or to relax positively some equation $q \in \mathcal{B}$ such that $\bar{x}_q > u_q^+$; or else it declares optimality if $u_B^- \leq \bar{x}_B \leq u_B^+$.

It remains to extend the rule for determining what new vertex is reached as we move along the edge. As before, a non-tight equation $j \in \mathcal{N}$ becomes tight when $(\bar{\pi}a_j - c_j) - \theta_q(\sigma_q a_j) = 0$, where $\sigma_q B = -e_q$. The difference now is that θ_q may be either increasing or decreasing from 0, while also the quantities $\bar{\pi}a_j - c_j$ may have either sign. As a result, the vertex is reached either when θ_q decreases to

$$\Theta_q^- = \max_{<0} \frac{\bar{\pi}a_j - c_j}{\sigma_q a_j},$$

or when θ_q increases to

$$\Theta_q^+ = \min_{>0} \frac{\bar{\pi}a_j - c_j}{\sigma_q a_j},$$

where the min is over all $j \in \mathcal{N}$ such that the ratios that are positive, and the max is over all $j \in \mathcal{N}$ such that the ratios that are negative. To define this ratio test more precisely (and in a way that will more clearly extend to the degenerate case), we write

$$\mathcal{N}^- = \{j : \bar{\pi}a_j - c_j < 0\} \quad \text{and} \quad \mathcal{N}^+ = \{j : \bar{\pi}a_j - c_j > 0\}.$$

Then a vertex is reached either when θ_q decreases to

$$\Theta_q^- = \max \left\{ \max_{\substack{j \in \mathcal{N}^- \\ \sigma_q a_j > 0}} \frac{\bar{\pi}a_j - c_j}{\sigma_q a_j}, \max_{\substack{j \in \mathcal{N}^+ \\ \sigma_q a_j < 0}} \frac{\bar{\pi}a_j - c_j}{\sigma_q a_j} \right\},$$

or when θ_q increases to

$$\Theta_q^+ = \min \left\{ \min_{\substack{j \in \mathcal{N}^- \\ \sigma_q a_j < 0}} \frac{\bar{\pi}a_j - c_j}{\sigma_q a_j}, \min_{\substack{j \in \mathcal{N}^+ \\ \sigma_q a_j > 0}} \frac{\bar{\pi}a_j - c_j}{\sigma_q a_j} \right\}.$$

The dual simplex method chooses some index $p \in \mathcal{N}$ at which the maximum is achieved (if equation q is being relaxed negatively) or at which the minimum is achieved (if equation q is being relaxed positively). Then, as before, p joins \mathcal{B} while q joins \mathcal{N} .

The iteration can also be finished as before, by changing one column of B and updating the values $\bar{s}_j = \bar{\pi}a_j - c_j$ by the previously given formula. If equation q has been relaxed negatively, then \bar{u}_q is set to u_q^- for the next iteration, and the vector $b - N\bar{u}_N$ is updated by adding $a_p \bar{u}_p$ and subtracting $a_q u_q^-$; similarly, if equation q has been relaxed positively, then \bar{u}_q is set to u_q^+ for the next iteration, and the vector $b - N\bar{u}_N$ is efficiently updated by adding $a_p \bar{u}_p$ and subtracting $a_q u_q^+$.

The primal problem associated with (\mathbf{D}) is precisely the bounded-variable linear program,

$$\begin{aligned} (\mathbf{P}) \quad & \text{Minimize} && cx \\ & \text{Subject to} && Ax = b \\ & && u^- \leq x \leq u^+ \end{aligned}$$

What we have derived above is a dual simplex method for this problem. Again, the dual method can use the same data structures and linear algebra routines as the standard primal simplex method for the associated primal LP.

2. The dual simplex method for (\mathbf{D}) extends naturally to permit some of the u_j^- to be $-\infty$, and some of the u_j^+ to be $+\infty$. Since the method will never choose a $q \in \mathcal{B}$ such that $\bar{x}_{Bq} < u_q^- = -\infty$, or such that $\bar{x}_{Bq} > u_q^+ = +\infty$, \bar{u}_q is always set to a finite value at the conclusion of an iteration. Thus, provided there are no infinities in \bar{u} at the first iteration, no infinities can be introduced into \bar{u} , and all of the formulas of the method carry through as before.

In more familiar terms, we can say that a basis matrix B gives a basic feasible solution in (\mathbf{D}) if

$$\begin{aligned} \bar{\pi}a_j - c_j &\geq 0 && \text{for each } u_j^- = -\infty, \\ \bar{\pi}a_j - c_j &\leq 0 && \text{for each } u_j^+ = +\infty. \end{aligned}$$

In the special case where all $u_j^- = 0$ and all $u_j^+ = +\infty$, this reduces to the notion of a basic feasible solution in (\mathbf{D}_0) . Indeed in this case (\mathbf{D}) becomes (\mathbf{D}_0) , and the algorithm of §1 reduces to the simpler algorithm of §0.

If a basic feasible solution is not known initially, one may be found by solving the auxiliary problem

$$(\mathbf{D}_f) \quad \text{Maximize} \quad -(\pi A - c)v^\pm$$

where

$$v_j^- = \begin{cases} -1 & \text{if } u_j^- = -\infty \\ 0 & \text{otherwise} \end{cases}, \quad v_j^+ = \begin{cases} +1 & \text{if } u_j^+ = +\infty \\ 0 & \text{otherwise} \end{cases};$$

this is simply the minimization of the “sum of infeasibilities” associated with π . Since there are no infinities in v^- or v^+ , all basic solutions for this problem are feasible, and hence it may be solved by the dual simplex method of §2 starting from *any* basic solution. If a basic solution having an objective value of 0 is found, then it can be used as an initial basic feasible solution for solving (\mathbf{D}) by our dual simplex method. Otherwise, no basic feasible solution exists.

With these extensions, the dual simplex method for (\mathbf{D}) accommodates any combination of upper and lower bounds in (\mathbf{P}) , and applies as well to a natural “phase I” problem for finding an initial dual feasible solution.

Various ideas for taking advantage of the phase I problem’s special form, as described by Wolfe for the primal simplex method, can also be applied to this dual method. At the beginning of an iteration we have, as in §1, a basis matrix B , a current iterate $\bar{\pi}$ satisfying $\bar{\pi}B = c_B$, and associated values \bar{u}_j which may or may not be finite. (It might seem that we should define values \bar{v}_j , since the objective is $-(\pi A - c)v^\pm$, but the relationship between v^\pm and u^\pm is so close that we can more clearly state this approach in terms of the latter.)

The equations for \bar{x} in phase I reduce to

$$B\bar{x}_B = b + \sum_{\bar{u}_j = -\infty} a_j - \sum_{\bar{u}_j = +\infty} a_j.$$

The selection of a basic equation to relax may also be simplified, by considering only relaxations in a feasible direction. That is, the method may choose an equation $q \in \mathcal{B}$ to be relaxed

$$\begin{aligned} &\text{along the negative edge, if } \bar{x}_q < 0 \text{ and } u_q^- > -\infty, \\ &\text{along the positive edge, if } \bar{x}_q > 0 \text{ and } u_q^+ < +\infty. \end{aligned}$$

If $\bar{x}_i \geq 0$ for all $i \in \mathcal{B} : u_i^- > -\infty$, and $\bar{x}_i \leq 0$ for all $i \in \mathcal{B} : u_i^+ < +\infty$, then the method can stop and declare that no feasible solution exists, even though some further reduction in the sum of infeasibilities may be possible.

Once $q \in \mathcal{B}$ is chosen, the step along the corresponding edge may be allowed to *pass through vertices*, as long as no new infeasibilities are created. To describe the appropriate ratio test, we further partition the nonbasic constraints $j \in \mathcal{N}^- \cup \mathcal{N}^+$ into those that will be infeasible if their slack changes sign,

$$\mathcal{F}_\infty^- = \{j \in \mathcal{N}^- : u_j^+ = +\infty\}, \quad \mathcal{F}_\infty^+ = \{j \in \mathcal{N}^+ : u_j^- = -\infty\},$$

and those that will be feasible if their slack changes sign:

$$\mathcal{F}_0^- = \{j \in \mathcal{N}^- : u_j^+ < +\infty\}, \quad \mathcal{F}_0^+ = \{j \in \mathcal{N}^+ : u_j^- > -\infty\}.$$

We then choose the either first constraint $p \in \mathcal{F}_\infty^- \cup \mathcal{F}_\infty^+$ that becomes tight, or the *last* constraint $p \in \mathcal{F}_0^- \cup \mathcal{F}_0^+$ that becomes tight—whichever becomes tight sooner as we move along the edge. More precisely, if equation q is being relaxed negatively, we let θ_q decrease to the greater of

$$\Theta_q^- = \max \left\{ \max_{\substack{j \in \mathcal{F}_\infty^- \\ \sigma_q a_j > 0}} \frac{\bar{\pi} a_j - c_j}{\sigma_q a_j}, \max_{\substack{j \in \mathcal{F}_\infty^+ \\ \sigma_q a_j < 0}} \frac{\bar{\pi} a_j - c_j}{\sigma_q a_j} \right\}$$

and

$$\Phi_q^- = \min \left\{ \min_{\substack{j \in \mathcal{F}_0^- \\ \sigma_q a_j > 0}} \frac{\bar{\pi} a_j - c_j}{\sigma_q a_j}, \min_{\substack{j \in \mathcal{F}_0^+ \\ \sigma_q a_j < 0}} \frac{\bar{\pi} a_j - c_j}{\sigma_q a_j} \right\};$$

if equation q is being relaxed positively, we let θ_q increase to the lesser of

$$\Theta_q^+ = \min \left\{ \min_{\substack{j \in \mathcal{F}_\infty^- \\ \sigma_q a_j < 0}} \frac{\bar{\pi} a_j - c_j}{\sigma_q a_j}, \min_{\substack{j \in \mathcal{F}_\infty^+ \\ \sigma_q a_j > 0}} \frac{\bar{\pi} a_j - c_j}{\sigma_q a_j} \right\}$$

and

$$\Phi_q^+ = \max \left\{ \max_{\substack{j \in \mathcal{F}_0^- \\ \sigma_q a_j < 0}} \frac{\bar{\pi} a_j - c_j}{\sigma_q a_j}, \max_{\substack{j \in \mathcal{F}_0^+ \\ \sigma_q a_j > 0}} \frac{\bar{\pi} a_j - c_j}{\sigma_q a_j} \right\}.$$

It can be shown that, at every iteration using this test, either the number of infeasibilities is reduced, or the number of infeasibilities stays the same and the phase I objective—the sum of infeasibilities—is reduced; this is sufficient to guarantee termination.

This phase I ratio test need not cost significantly more than the test described in §1. The updates to the slacks \bar{s}_j are the same, while the updates to the linear system for \bar{x} are somewhat different, due to the special nature of the phase I objective and the possibility that many slacks may change sign. For each constraint $j \in \mathcal{F}_0^- \cup \mathcal{F}_0^+$ that has gone from infeasible to feasible as a result of the iteration, we must update the right-hand side of the system for \bar{x} by adding a_j (if we had $\bar{u}_j = -\infty$) or subtracting a_j (if we had $\bar{u}_j = +\infty$).

3. When $\bar{\pi}a_j = c_j$ for some nonbasic indices $j \in \mathcal{N}$, the current basic solution is degenerate, and we must refine several steps of the algorithm.

At the initial basic feasible solution, we choose the “current slopes” \bar{u}_j corresponding to the degenerate equations as

$$\bar{u}_j = \left\{ \begin{array}{ll} u_j^+ & \text{if } u_j^- = -\infty \\ u_j^- & \text{if } u_j^+ = +\infty \\ u_j^+ \text{ or } u_j^- & \text{otherwise} \end{array} \right\} \text{ for each } \bar{\pi}a_j = c_j, j \in \mathcal{N}.$$

Where both u_j^+ and u_j^- are finite, \bar{u}_j may be chosen as either one, by any arbitrary heuristic.⁶ Different choices for these components of \bar{u}_N lead to different solutions of $B\bar{x} = b - N\bar{u}_N$. Nevertheless, if $u_B^- \leq \bar{x}_B \leq u_B^+$ then the current basic solution is optimal, regardless of what choice was made.

If $\bar{x}_q < u_q^-$ or $\bar{x}_q > u_q^+$ for some $q \in \mathcal{B}$, then the current basic solution may not be optimal, and the basic equation corresponding to index q may be relaxed as before. The previous definition of the sets \mathcal{N}^- and \mathcal{N}^+ is no longer sufficient, however, because it does not say in which of these sets to place an index $j \in \mathcal{N}$ such that $\bar{\pi}a_j - c_j = 0$. We must instead define them as

$$\mathcal{N}^- = \{j : \bar{u}_j = u_j^-\} \quad \text{and} \quad \mathcal{N}^+ = \{j : \bar{u}_j = u_j^+\}.$$

Thus if $\bar{\pi}a_j - c_j = 0$ and we chose $\bar{u}_j = u_j^-$, we must be consistent by placing j in \mathcal{N}^- ; or if we chose $\bar{u}_j = u_j^+$, we must be consistent by placing j in \mathcal{N}^+ . For any j such that $\bar{\pi}a_j - c_j < 0$ or $\bar{\pi}a_j - c_j > 0$, we are back in the nondegenerate case, and the above definition puts j into the same set as previously.

Given these refined definitions of \mathcal{N}^- and \mathcal{N}^+ , the statement of the ratio test itself is unchanged, as is the rest of the iteration. In the updating at the end, we must take care to also set \bar{u}_q in a consistent way: to u_q^- if we tried to relax equation q negatively, or to u_q^+ if we tried to relax equation q positively.

Since degeneracy involves $\bar{\pi}a_j - c_j = 0$ for certain j , the ratio test may produce $\Theta_q^- = 0$ or $\Theta_q^+ = 0$, with the result being a null step that changes the basis but not the basic solution $\bar{\pi}$ or the slacks \bar{s} . We conclude by looking at a generalized phase II ratio test that may be less susceptible to null steps.

⁶The case of $u_j^- = -\infty$ and $u_j^+ = +\infty$, corresponding to a free variable in the primal LP, has to be handled specially; but it poses no serious obstacles to implementation.

4. In the dual simplex method as we have described it so far, an iteration moves from the current vertex, along an edge, to the vertex at the other end of that edge. When the method is applied to (\mathbf{D}_0) , an iteration must stop at the latter vertex, since otherwise the next iterate would be infeasible. When applied to (\mathbf{D}) , however, the iterate may move through one or more vertices, since the edge may be feasible on both sides of a vertex.

In geometric terms, the objective function along an edge is linear between adjacent vertices, and concave piecewise-linear overall. As a result, each time that the iterate passes through a vertex, the rate of increase in the objective becomes less favorable. Eventually, past some vertex (unless the objective is unbounded), the rate of increase becomes negative, and the iteration must stop at that vertex and add the index of the corresponding tight constraint to \mathcal{B} .

In algebraic terms, if both u_j^- and u_j^+ are finite for some $j \in \mathcal{N}$, then the slack $\bar{\pi}a_j - c_j$ may be allowed to change sign in the course of an iteration. This generalization affects the dual simplex method only in the ratio test to determine which equation $p \in \mathcal{N}$ will become tight and join the set of basic equations. Specifically, p need no longer correspond to the maximum ratio (if a basic equation is being relaxed positively) or to the minimum ratio (if a basic equation is being relaxed negatively).

To make this observation precise, suppose that a dual simplex iteration is increasing θ_q so as to move along the positive edge associated with relaxing constraint q . The slack on each $j \in \mathcal{N}$ is then changing (as seen originally in §0) to

$$\pi(\theta_q)a_j - c_j = (\bar{\pi}a_j - c_j) - \theta_q(\sigma_q a_j).$$

Thus the slack in the j th constraint is a linear function of θ , with a slope of $\sigma_q a_j$. It follows that the objective term $-(\pi(\theta_q)a_j - c_j)u_j^\pm$ is also a linear function of θ , with slope $(\sigma_q a_j)\bar{u}_j$, so long as $\pi(\theta_q)a_j - c_j$ does not change sign.

Suppose next that we have determined Θ_q^+ such that $\pi(\Theta_q^+)a_p - c_p = 0$. If we continue past this vertex to points where $\theta > \Theta_q^+$, then the slope of $(\pi a_p - c_p)u_p^\pm$ changes. There are two cases:

- $\sigma_q a_p < 0$: the slack changes from negative to positive; the slope changes from $(\sigma_q a_p)u_p^-$ to $(\sigma_q a_p)u_p^+$.
- $\sigma_q a_p > 0$: the slack changes from positive to negative; the slope changes from $(\sigma_q a_p)u_p^+$ to $(\sigma_q a_p)u_p^-$.

Either way, the slope decreases by $|\sigma_q a_p|(u_p^+ - u_p^-) \geq 0$. The effect is to make the objective increase at a lower rate as we move further along the edge, or to make it start decreasing.

We know the initial slope of the objective along the positive edge; it is given (as shown in §1) by $\bar{x}_q - u_q^+$. Thus we can consider moving through a series of vertices corresponding to constraints $j \in \mathcal{N}$, reducing the slope by $|\sigma_q a_j|(u_j^+ - u_j^-)$ at each one, and stopping at the vertex where the slope first becomes nonpositive. To state this precisely, we denote by \mathcal{Q}^+ the set of all constraints $j \in \mathcal{N}$ whose slacks may change sign if we move far enough along the positive edge:

$$\mathcal{Q}^+ = \{j \in \mathcal{N}^- : \sigma_q a_j < 0\} \cup \{j \in \mathcal{N}^+ : \sigma_q a_j > 0\}.$$

We also observe that the values from the dual simplex ratio test,

$$\frac{\bar{\pi}a_j - c_j}{\sigma_q a_j}, \quad j \in \mathcal{Q}^+,$$

are the values of θ_q at which the slacks do change sign. Thus we can consider these values in increasing order, adjusting the slope of the objective until it ceases to be positive. The logic is as follows:

Initialize $d_q \leftarrow \bar{x}_q - u_q^+$, and \mathcal{Q}^+ as above.
Repeat while \mathcal{Q}^+ is not empty and $d_q > 0$:
 Choose $p = \arg \min_{j \in \mathcal{Q}^+} (\bar{\pi}a_j - c_j) / \sigma_q a_j$.
 Let $d_q \leftarrow d_q - |\sigma_q a_p| (u_p^+ - u_p^-)$.
 Let $\mathcal{Q}^+ \leftarrow \mathcal{Q}^+ \setminus \{p\}$.
Let $\Theta_q^+ \leftarrow (\bar{\pi}a_p - c_p) / \sigma_q a_p$.

This loop replaces the ratio test of §1 for determining q . If it encounters a $u_p^- = -\infty$ or $u_p^+ = +\infty$, then d_q becomes infinitely negative, and the loop is terminated; any further movement along the edge will be outside the feasible region. The computational work need not be substantially greater than for the simpler ratio test in §1, however. Infinities in u^- and u^+ can be handled specially, and the successive ratios can be retrieved from an efficient data structure such as a binary heap.

Once q is determined by the generalized ratio test, the iteration is concluded as before. The only extra computation occurs in updating the right-hand side of $B\bar{x}_B = b - \sum_{j \in \mathcal{N}} a_j \bar{u}_j$ to account for slacks that change sign. This work can be done as part of the above loop, by starting with $b' \leftarrow b - \sum_{j \in \mathcal{N}} a_j \bar{u}_j$ and updating in each pass:

If $\sigma_q a_p > 0$ **then** $b' \leftarrow b' + a_p u_p^- - a_p u_p^+$.
If $\sigma_q a_p < 0$ **then** $b' \leftarrow b' + a_p u_p^+ - a_p u_p^-$.

More concisely, $b' \leftarrow b' - a_p \operatorname{sign}(\sigma_q a_p) (u_p^+ - u_p^-)$.

The generalized ratio test is intuitively appealing, because in effect it carries out a series of ordinary dual simplex iterations, at not much more than the cost of one such iteration. There is some minimum number of iterations, however, separating the starting basis from an optimal one. If the dual simplex method comes close to this minimum, then no generalization of the ratio test can give much improvement. If there are many infinities in u^- and u^+ , moreover, or even if $u_j^+ - u_j^-$ tends to be large, then the generalized ratio test may end up stopping at the first vertex much of the time anyway.

Studies of other piecewise-linear simplex methods suggest that the greatest advantage of the generalized ratio test may be in dealing with degeneracy, because it can succeed in taking a positive step even when some of the ratios are zero. In brief, the initial passes through the loop, where $(\bar{\pi}a_p - c_p) / \sigma_p a_j = 0$, can be regarded as “correcting” the value of d_q for the cases where we have “guessed wrong” about the slope associated with $p \in \mathcal{N}$: either $\sigma_q a_p < 0$ but $\bar{u}_p = u_p^-$, or $\sigma_q a_p > 0$ but $\bar{u}_p = u_p^+$. As a result, the loop may determine that some positive step can be taken along the edge, whereas the ordinary ratio test would have taken a degenerate step.

(Experience also suggests that, if the loop finds that a positive step cannot be taken, then the \bar{u}_j , $j \in \mathcal{N}$ should be left at their previous values; changing these values tends only to produce a greater number of degenerate steps.)

The generalized ratio test can be applied just as well to the phase I problem, (\mathbf{D}_f) . However, there is as yet no computational evidence to suggest that any approach is better for phase I than the Wolfe's simpler ratio test.