Closing the gap in pivot methods for linear programming

Venkat Narayan
Venkat.Narayan@outlook.com

Abstract

We propose pivot methods that solve linear programs by trying to close the duality gap from both ends. The first method maintains a set $\mathcal{B}$ of at most three bases, each of a different type, in each iteration: a primal feasible basis B^p, a dual feasible basis B^d and a primal-and-dual infeasible basis B^i. From each $B \in \mathcal{B}$, it evaluates the primal and dual feasibility of all primal and dual pivots to maximize lower bound $\underline{z}$ and minimize upper bound $\overline{z}$ on the optimal objective value z^*; i.e., it explores the adjacency neighborhood of the primal-dual vertex pair corresponding to B in the underlying primal and dual hyperplane arrangements to minimize duality gap $\overline{z} - \underline{z}$. Among all primal feasible bases with (non-decreasing) objective $z = \underline{z}$ adjacent to any $B \in \mathcal{B}$, it selects the least infeasible B'^p to update B^p. Likewise, it updates B^d to the least infeasible B'^d with (non-increasing) $z = \overline{z}$, and B^i to the least infeasible B'^i with non-monotone $z \in [\underline{z},\ \overline{z}]$. Bases B'^p, B'^d and B'^i may result from any $B \in \mathcal{B}$ enabling different interconnected, non-adjacent, criss-cross paths to approach a terminal basis. Ignoring B^i and primal-and-dual infeasible pivots, a variant maintains and updates at most two bases B^p and B^d. Restricting both methods to work with one customary basis yields criss-cross variants that traverse a non-monotone path visiting intermingled least infeasible bases (a) B'^p, B'^d or B'^i, and (b) B'^p or B'^d in any order. To implement them, we outline primal and dual feasibility tests that complement the simplex minimum ratio test. Resolving the finiteness and complexity of these methods is an open problem.

1 Introduction

The linear programming problem [10] is to optimize a linear objective function subject to linear constraints on non-negative decision variables. Consider the following linear program (LP):

$$\max\{\hat{c}^{\mathsf{T}}\hat{x} : \ \hat{x} \in \mathcal{P}\} \text{ where } \mathcal{P} = \{\hat{x} : \ \hat{A}\hat{x} \le \hat{b}, \ \hat{x} \ge 0\},$$
$$\hat{b} \in \mathcal{R}^m, \ \hat{c}, \ \hat{x} \in \mathcal{R}^n \text{ and } \hat{A} \in \mathcal{R}^{m\times n} \text{ has rank } m.$$

The dual LP corresponding to the above primal LP may be written as follows where $\hat{y} \in \mathcal{R}^m$:

$$\min\{\hat{b}^{\mathsf{T}}\hat{y} : \ \hat{y} \in \mathcal{D}\} \text{ where } \mathcal{D} = \{\hat{y} : \ \hat{A}^{\mathsf{T}}\hat{y} \ge \hat{c}, \ \hat{y} \ge 0\}.$$

Using primal slack variables $\hat{u} \in \mathcal{R}^m, \ \hat{u} \ge 0$, we may write the primal constraints as $\hat{A}\hat{x}+\hat{u} = \hat{b}$. Using dual excess variables $\hat{v} \in \mathcal{R}^n, \ \hat{v} \ge 0$, we may write the dual constraints as $\hat{A}^{\mathsf{T}}\hat{y} - \hat{v} = \hat{c}$.

If $\mathcal{P} = \emptyset$ (resp. $\mathcal{D} = \emptyset$), we say the primal (resp. dual) LP is inconsistent, else it is consistent.

We say $\hat{x} \notin \mathcal{P}$ is primal infeasible, $\hat{x} \in \mathcal{P}$ is primal feasible, $\hat{y} \notin \mathcal{D}$ is dual infeasible and $\hat{y} \in \mathcal{D}$ is dual feasible. If $\hat{x} \in \mathcal{P}$ and $\hat{y} \in \mathcal{D}$, then $\hat{x}^{\mathsf{T}}\hat{c} \le \hat{x}^{\mathsf{T}}\hat{A}^{\mathsf{T}}\hat{y} \le \hat{b}^{\mathsf{T}}\hat{y}$; i.e., any $\hat{x} \in \mathcal{P}$ (resp. $\hat{y} \in \mathcal{D}$) yields a lower (resp. an upper) bound on the dual (resp. primal) objective value. This is *weak duality*. The difference $\hat{b}^{\mathsf{T}}\hat{y} - \hat{c}^{\mathsf{T}}\hat{x} = (\hat{x}^{\mathsf{T}}\hat{A}^{\mathsf{T}} + \hat{u}^{\mathsf{T}})\hat{y} - \hat{x}^{\mathsf{T}}(\hat{A}^{\mathsf{T}}\hat{y} - \hat{v}) = \hat{u}^{\mathsf{T}}\hat{y} + \hat{v}^{\mathsf{T}}\hat{x} \ge 0$ is the *duality gap* with respect to $\hat{x}$ and $\hat{y}$. The goal is to find a feasible primal-dual $(\hat{x}, \ \hat{y})$ pair so the gap is zero, or certify that no such pair exists. By *strong duality*, a primal optimum $\hat{x}^*$ exists if and only if a dual optimum $\hat{y}^*$ exists with $\hat{c}^{\mathsf{T}}\hat{x}^* = \hat{b}^{\mathsf{T}}\hat{y}^*$. Equivalently, $(\hat{x}^*, \ \hat{v}^*)$ and $(\hat{y}^*, \ \hat{u}^*)$ are complementary: $(\hat{v}^*)^{\mathsf{T}}\hat{x}^* = 0$ and $(\hat{u}^*)^{\mathsf{T}}\hat{y}^* = 0$. This is *complementary slackness* or *complementarity*. Such a pair exists if and only if the primal and dual LPs are consistent. Thus, solution pair $(\hat{x}, \ \hat{y})$ is optimal if and only if $\hat{x}$ is primal feasible, $\hat{y}$ is dual feasible, and $(\hat{x}, \ \hat{v})$ and $(\hat{y}, \ \hat{u})$ are complementary [16, 32]:

$$\hat{A}\hat{x} + \hat{u} = \hat{b} \quad (1a) \qquad \hat{A}^{\mathsf{T}}\hat{y} - \hat{v} = \hat{c} \quad (1c) \qquad \hat{v}^{\mathsf{T}}\hat{x} = 0 \quad (1e)$$
$$\hat{x}, \ \hat{u} \ge 0 \quad (1b) \qquad \hat{y}, \ \hat{v} \ge 0 \quad (1d) \qquad \hat{u}^{\mathsf{T}}\hat{y} = 0 \quad (1f)$$

The system of primal (resp. dual) hyperplanes $\hat{x} = 0$ and $\hat{A}\hat{x} = \hat{b}$, n.b., $\hat{u} = 0$ (resp. $\hat{y} = 0$ and $\hat{A}^{\mathsf{T}}\hat{y} = \hat{c}$, n.b., $\hat{v} = 0$) partitions $\mathcal{R}^n$ (resp. $\mathcal{R}^m$) into connected convex polyhedral cells or regions [40]. This primal (resp. dual) hyperplane arrangement forms 1-skeleton graph $\mathcal{A}_P$ (resp. $\mathcal{A}_D$) with vertices at which n (resp. m) or more hyperplanes intersect and edges at which $n-1$ (resp. $m-1$) hyperplanes meet to link vertex pairs. Convex polyhedron $\mathcal{P}$ (resp. $\mathcal{D}$) is its possibly empty feasible region. Edges of $\mathcal{P}$ (resp. $\mathcal{D}$) connect its extreme points to form graph $\mathcal{E}_P \subseteq \mathcal{A}_P$ (resp. $\mathcal{E}_D \subseteq \mathcal{A}_D$).

Intense work on LPs since the 1940s has led to influential progress in theory, models, algorithms, computations and applications [35, 3, 27]. Rich expressive power and efficient solvability have made it a cornerstone of mathematical programming. But answers to some fundamental questions remain elusive [17, 38]: Does LP admit a strongly polynomial algorithm in the arithmetic model of computation? Does LP admit a polynomial algorithm in the real number model of computation, i.e., one whose count of unit cost operations on real numbers is bounded by a polynomial $P(n, \ m)$? If such methods exist, a further challenge is to design implementations that are efficient in practice.

A strongly polynomial method [17] would use basic arithmetic operations $(+, \ -, \ \times, \ \div, \ <)$ with operation count (resp. space use) bounded by $P(n, \ m)$ (resp. $P(n, \ m, \ \mathcal{L})$), establishing the inherent complexity of LP to be independent of bit size $\mathcal{L} = \mathcal{L}(\hat{A}) + \mathcal{L}(\hat{b}) + \mathcal{L}(\hat{c})$. LP admits (weakly) polynomial ellipsoid and interior point methods [24, 22], among others, with both measures bounded by $P(n,\ m,\ \mathcal{L})$. The former has unbounded complexity in the real number model [46].

Many specialized LPs admit strongly polynomial methods, e.g., fixed dimensional LPs [30], minimum cost flows (MCF) [42], combinatorial LPs via polynomial methods bounded by $P(n,\ m,\ \mathcal{L}(\hat{A}))$ [43], feasibility of LPs with at most two variables per inequality (TVPI) [29] or homogeneous constraints [7], deformed products [2] and certain Markovian decision problems (MDP) [47]. Their existence is open for others like the generalized MCF, its restriction optimizing TVPI, and MDP.

Pivot methods are natural candidates in the search for a strongly polynomial algorithm for LP. They use local search based strongly polynomial linear algebraic operations to walk a basis-pivot graph $\mathcal{G}$. Bases are invertible $m \times m$ submatrices of constraint matrix A. Pivots connect each basis to its neighbors. Usually, pivots are adjacency-based, i.e., they swap a column i in basis B with a $j \notin B$, thus traversing $\mathcal{A}_P$ and $\mathcal{A}_D$. A local optimum is globally optimal as LP optimizes a linear function over a convex polyhedron. In general, much of this work is characterized by geometrically or combinatorially inspired pivot rules with a low complexity iteration to select a neighboring basis [45, 14]. Many of them have exponential or subexponential worst case complexity [25, 1, 13].

The simplex framework [10] has preoccupied the search for a strongly polynomial pivot method in part due to its surprising effectiveness on real-world instances and unmatched dominance till the 1980s. Variants with polynomial average case [4], polynomial smoothed [39], randomized polynomial [23] and subexponential [20] complexity partially explain its good observed behavior.

A simplex path is a monotone, possibly degenerate walk on $\mathcal{E}_P$ (or resp. $\mathcal{E}_D$) using primal (resp. dual) feasible bases-simplex pivots in $\mathcal{G}$. For it to be polynomial, the diameter Δ_n^f of polyhedron $\mathcal{P}$ in $\mathcal{R}^n$, say, with f facets should be bounded by $P(n,\ f)$ non-degenerate, possibly non-monotone pivots; Δ_n^f has linear lower bounds and subexponential upper bounds [34, 21]. Simplex paths need at least Δ_n^f pivots in the worst case. The conjecture that Δ_n^f is polynomially bounded is open.

A pivot rule in [6] uses two adjacent bases that differ in one column: a primal feasible B^p and dual feasible B^d. It updates B^p (or resp. B^d) via a primal (resp. dual) simplex pivot with $B^d \backslash B^p$ as entering (resp. leaving) column e (resp. l). Subsequently, it uses l (resp. e) from the previous primal (resp. dual) simplex pivot as l (resp. e) in the next dual (resp. primal) simplex pivot to close the duality gap at both ends. For an exponential example and related methods, see [12].

Some methods walk monotone, possibly non-adjacent extreme points of $\mathcal{P}$ via block pivots [9, 36] or shortcuts through its interior [2], or traverse adjacent faces of different dimensions [5].

Criss-cross methods [48] can start from any basis and subsequently walk any path in $\mathcal{G}$ to solve the LP in one phase with non-monotone objective z. They may visit bases of any type to freely walk any possibly degenerate path in $\mathcal{A}_P$ and $\mathcal{A}_D$. For a review of criss-cross and related methods, see [14]. The self-dual parametric method [11] has exponential worst case [31] and polynomial average case [37] complexity. Others include combinatorial exponential [44, 33] and randomized subexponential [28] pivot rules, and a phase I simplex equivalent [26, 18] for linear feasibility. The diameter of a hyperplane arrangement is $\mathcal{O}(mn)$. A path of at most $m + n$ pivots exists in $\mathcal{G}$ from any basis to a terminal one [15]. It is intriguing to consider how such a short path may be realized.

Section 2 reviews bases, dictionaries and pivots. Based on a primal-dual perspective of $\mathcal{G}$, section 3 proposes pivot methods that suspend practicality to close the duality gap at both ends at the cost of more expensive iterations. Given bases $\mathcal{B}$ of different types, they evaluate the primal and dual feasibility of all pivots adjacent to each $B \in \mathcal{B}$ to maximize lower bound $\underline{z}$ and minimize upper bound $\overline{z}$ on the optimal objective value z^*. Using least infeasible pivots with $z \in [\underline{z},\ \overline{z}]$, they update $\mathcal{B}$ to traverse interconnected, non-adjacent, criss-cross paths. To implement them, section 4 outlines primal and dual feasibility tests that complement the simplex minimum ratio test.

2 Bases, Dictionaries and Pivots

Let $x = \hat{x} \cup \hat{u}$, $A' = \hat{A} \cup I_m$ and $c' = \binom{\hat{c}}{0} \in \mathcal{R}^{n+m}$. Consider a partition of primal variables x to *non-basic* (independent) variables x_N and *basic* (dependent) variables x_B where *basis* B, $|B| = m$, is nonsingular, i.e., A'_B is linearly independent. Write the primal LP as $\max\{c_B'^{\mathsf{T}} x_B + c_N'^{\mathsf{T}} x_N :\ A'_B x_B + A'_N x_N = \hat{b},\ x \geq 0\}$. Setting $x_N = 0$ and solving the system of equations $A'_B x_B = \hat{b}$ yields the primal *dictionary* (2) [41, 8] with $x_B = b = A_B'^{-1}\hat{b},\ D = -A_B'^{-1} A'_N,\ z = c_B'^{\mathsf{T}} b$ and $c_N = c_N'^{\mathsf{T}} + c_B'^{\mathsf{T}} D$:

$$x_i = b_i + \sum_{j \in N} D_{ij} x_j \qquad \forall i \in B \tag{2a}$$

$$z_P = z + \sum_{j \in N} c_j x_j. \tag{2b}$$

$$y_j = -c_j - \sum_{i \in B} D_{ij} y_i \qquad \forall j \in N \tag{3a}$$

$$-z_D = -z - \sum_{i \in B} b_i y_i. \tag{3b}$$

Let dual variables $y = \hat{y} \cup \hat{v}$. Dual non-basic (resp. basic) variables y_B (resp. y_N) are the complements of primal basic (resp. non-basic) variables x_B (resp. x_N). Dual dictionary (3) with $y_N = -c_N$ and $y_B = 0$ is a negative transpose of (2). Basis B represents a complementary basic solution pair $(x,\ y)$, i.e., $x^{\mathsf{T}} y = 0$, that satisfies all optimality conditions (1) except primal (1b) and dual (1d) non-negativity. The active hyperplanes implied by $x_N = 0$ (resp. $y_B = 0$) intersect at vertex $\hat{x} \in \mathcal{A}_P$ (resp. $\hat{y} \in \mathcal{A}_D$). If x_B and y_N are non-negative, then B would be optimal.

Define $L(B)$ (resp. $E(B)$) as the set of primal (resp. dual) infeasible variables for basis B:

$$L(B) = \{i \in B :\ b_i < 0\} \quad \text{and} \quad E(B) = \{j \in N :\ c_j > 0\}.$$

If $L(B) = \emptyset$ (resp. $E(B) = \emptyset$), then B is primal (resp. dual) feasible, i.e., $\hat{x} \in \mathcal{P}$ (resp. $\hat{y} \in \mathcal{D}$) is an extreme point of $\mathcal{P}$ (resp. $\mathcal{D}$), else B is primal (resp. dual) infeasible, i.e., $\hat{x} \notin \mathcal{P}$ (resp. $\hat{y} \notin \mathcal{D}$) is an infeasible vertex of $\mathcal{A}_P$ (resp. $\mathcal{A}_D$). Based on $L(B)$ and $E(B)$, there are four primal-dual basis types. If B satisfies any termination condition below, then B is *terminal*, else it is non-terminal.

$$\text{If } L(B) \cup E(B) = \emptyset,\ \text{then basis } B \text{ is optimal.} \tag{4a}$$

$$\text{If } \exists i \in L(B) \text{ with } D_{ij} \leq 0,\ \forall j \in N,\ \text{the LP is primal inconsistent.} \tag{4b}$$

$$\text{If } \exists j \in E(B) \text{ with } D_{ij} \geq 0,\ \forall i \in B,\ \text{the LP is dual inconsistent.} \tag{4c}$$

If B is non-terminal, we may try another and repeat until termination. To update B to a neighboring basis, many pivot methods use the so-called dual pivots in $\mathcal{C}_d(B)$ (and/or resp. primal pivots in $\mathcal{C}_p(B)$) below that result from primal (resp. dual) infeasible variables $L(B)$ (resp. $E(B)$):

$$\mathcal{C}_d(B) = \{(i,\ j):\ i \in L(B),\ j \in N,\ D_{ij} > 0\} \ \text{ and } \ \mathcal{C}_p(B) = \{(i,\ j):\ j \in E(B),\ i \in B,\ D_{ij} < 0\}.$$

In the sequel, when referencing a generic basis, we drop B and simply write $L,\ E,\ \mathcal{C}_p,\ \mathcal{C}_d$, etc.

The bases reachable via pivots $\mathcal{C} = \mathcal{C}_p \cup \mathcal{C}_d$ form the adjacency neighborhood of B in $\mathcal{G}$. For any pivot $a = (i,\ j) \in \mathcal{C}$, we may replace B with adjacent basis $B' = B \backslash i \cup j$. Swapping primal (resp.

dual) variable x_i in x_B with x_j in x_N (resp. y_j in y_N with y_i in y_B), i.e., replacing active primal (resp. dual) hyperplane $x_j = 0$ with $x_i = 0$ (resp. $y_i = 0$ with $y_j = 0$), represents a primal-dual transition from the $(x,\ y)$ pair of B to the $(x',\ y')$ pair of B'. In $\mathcal{A}_P$ (resp. $\mathcal{A}_D$), $\hat{x}'$ (resp. $\hat{y}'$) corresponds to an adjacent vertex along edge $[\hat{x},\ \hat{x}']$ (resp. $[\hat{y},\ \hat{y}']$), or to a degenerate vertex $\hat{x}' = \hat{x}$ (resp. $\hat{y}' = \hat{y}$) with different basis representations if $b_i = 0$ (resp. $c_j = 0$). For pivot $a = (i,\ j)$, let

$$b^a_j = -b_i/D_{ij} \quad (5a) \qquad c^a_i = c_j/D_{ij} \quad (5e)$$

$$b^a_p = b_p + b^a_j D_{pj}, \quad \forall p \in B\backslash i \quad (5b) \qquad c^a_q = c_q - c^a_i D_{iq}, \quad \forall q \in N\backslash j \quad (5f)$$

$$w_a(B) = b_i c_j / D_{ij} \quad (5c) \qquad z_a(B) = z - w_a(B) \quad (5g)$$

$$L_a(B) = \{p \in B\backslash i \cup j :\ b^a_p < 0\} \quad (5d) \qquad E_a(B) = \{q \in N\backslash j \cup i :\ c^a_q > 0\} \quad (5h)$$

If pivot a is selected, then $b = b^a,\ c = c^a,\ z = z_a(B),\ L(B') = L_a(B)$ and $E(B') = E_a(B)$ for basis B'. If $L_a = \varnothing$ (resp. $E_a = \varnothing$), pivot a, i.e., basis B', is primal (resp. dual) feasible, else it is primal (resp. dual) infeasible. Figure 1 classifies primal and dual pivots based on primal and dual feasibility. Parts of an arrangement edge may correspond to different pivot types. Primal inbound (resp. simplex) pivots are primal feasible pivots from a primal infeasible (resp. feasible) vertex $\hat{x} \notin \mathcal{P}$ (resp. $\hat{x} \in \mathcal{P}$) to $\hat{x}' \in \mathcal{P}$; $[\hat{x},\ \hat{x}']$ is an edge of $\mathcal{A}_P \backslash \mathcal{E}_P$ (resp. $\mathcal{E}_P$ if non-degenerate). Primal infeasible pivots result in $\hat{x}' \notin \mathcal{P}$ from any $\hat{x} \in \mathcal{A}_P$. Dual inbound, dual simplex and dual infeasible pivots are their equivalents with regard to $\mathcal{D}$, $\mathcal{A}_D$ and $\mathcal{E}_D$. Figure 2 is a primal-dual classification of bases and pivots in $\mathcal{G}$. Not all basis and pivot types may exist in $\mathcal{G}$, e.g., if $\mathcal{P} = \varnothing$ or $\mathcal{D} = \varnothing$.

Each non-terminal basis has a non-decreasing primal (resp. non-increasing dual) simplex pivot with $z_a \geq \underline{z}$ (resp. $z_a \leq \overline{z}$) if it is primal (resp. dual) feasible, else it may have a primal (resp. dual) inbound pivot with $z_a \geq \underline{z}$ (resp. $z_a \leq \overline{z}$). The primal (resp. dual) simplex method only considers the former to preserve primal (resp. dual) feasibility and monotonically improve $\underline{z}$ (resp. $\overline{z}$).

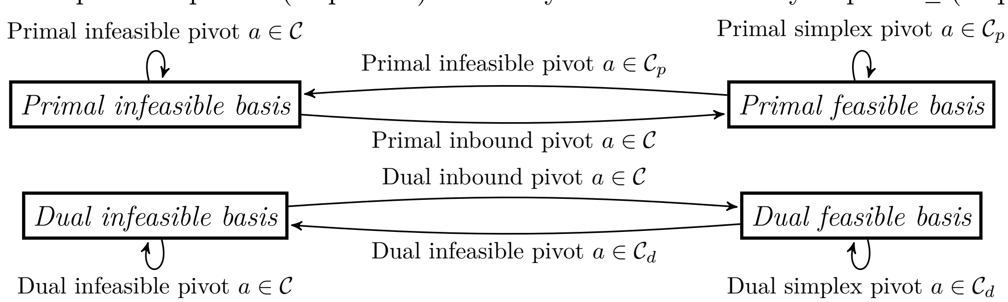

Figure 1: Types of primal and dual pivots with regard to primal feasibility and dual feasibility

3 Proposed Methods

We propose self-dual scale-invariant methods that use the rich structure of $\mathcal{G}$ alluded to in figure 2 to enable diverse ways to approach a terminal basis. They may be initialized with a primal (and/or resp. dual) feasible basis; if one is not available, we apply them on a primal (resp. dual) phase I LP with artificial variables (resp. constraints) to find one or establish primal (resp. dual) inconsistency.

3.1 Three Basis Method

Let $R = \{p,\ d,\ i\}$ be the allowed basis types. The three basis method maintains a set $\mathcal{B} \subseteq \cup_{r \in R} B^r$ of at most three bases, each of a different type, in each iteration: a primal feasible basis B^p, a dual feasible basis B^d and a primal-and-dual infeasible basis B^i. Initialize $\varnothing \neq \mathcal{B} \subseteq \{B^p,\ B^d\}$.

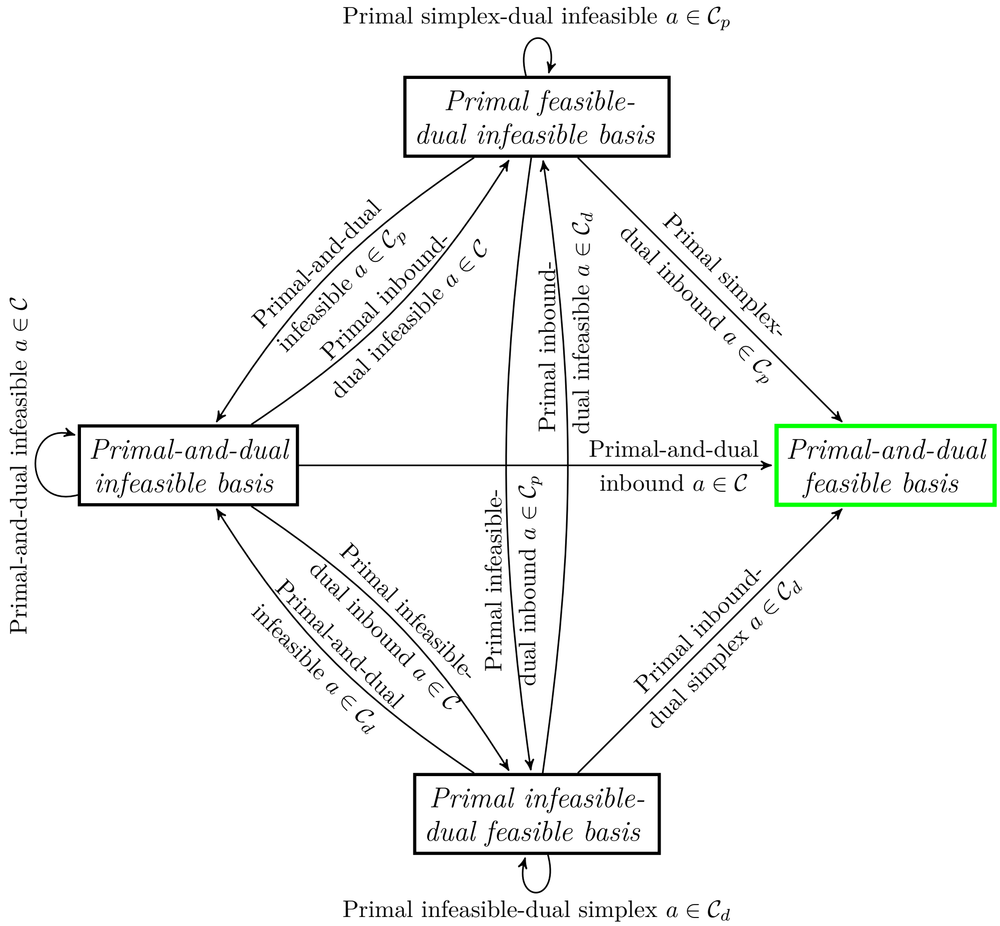

Figure 2: Primal-dual classification of bases and pivots

Step 0 initializes $\underline{z}$ (resp. $\overline{z}$) based on the primal (resp. dual) feasibility of initial bases $B \in \mathcal{B}$.

Step 1 examines if any $B \in \mathcal{B}$ is terminal by checking if it meets any termination condition (4).

Step 2 explores the neighborhood of each $B \in \mathcal{B}$ in $\mathcal{G}$, i.e., the neighbors of its corresponding $\hat{x}$ in $\mathcal{A}_P$ (resp. $\hat{y}$ in $\mathcal{A}_D$), to improve $\underline{z}$ (resp. $\overline{z}$) using primal (resp. dual) simplex or inbound pivots.

Step 3 uses $\underline{z}$ and $\overline{z}$ as pincers to guide the search by only considering pivots $a \in \mathcal{C}(B^o)$ with $z_a(B^o) \in [\underline{z},\ \overline{z}]$ from each basis $B^o \in \mathcal{B}$. The same pivot a may exist in any $\mathcal{C}(B^o)$ with different primal-dual consequences. Based on $L_a(B^o)$ and $E_a(B^o)$, it classifies each unpruned pivot a along with its originating basis B^o into three sets: primal feasible $\mathcal{W}^p(\mathcal{B})$, dual feasible $\mathcal{W}^d(\mathcal{B})$ and primal-and-dual infeasible $\mathcal{W}^i(\mathcal{B})$ where $|\mathcal{W}^p(\mathcal{B}) \cap \mathcal{W}^d(\mathcal{B})| \geq 0$. Each pivot-and-originating basis $(a,\ o)$ in (a) $\mathcal{W}^p(\mathcal{B})$ has $z_a(B^o) = \underline{z}$, (b) $\mathcal{W}^d(\mathcal{B})$ has $z_a(B^o) = \overline{z}$, and (c) $\mathcal{W}^i(\mathcal{B})$ has $z_a(B^o) \in [\underline{z},\ \overline{z}]$.

It is unclear which candidate $(a,\ o)$ in $\mathcal{W}^r(\mathcal{B}),\ r \in R$, should be selected to minimize the duality gap in the next iteration. Let $\pi(B) = f(L(B),\ E(B)) \geq 0$ be any non-decreasing function measuring the primal and dual infeasibility of B with $\pi(B) = 0$ if and only if B is optimal (4a). To look ahead and evaluate the primal-dual consequences of all $(a,\ o)$ in $\mathcal{W}^r(\mathcal{B})$ on a common basis, let $\pi_{ao} = f(L_a(B^o),\ E_a(B^o)) = \pi(B^o \backslash i \cup j)$, i.e., the infeasibility of pivot $a = (i,\ j)$ in B^o equals that of resulting basis $B^o \backslash i \cup j$. Total or maximum infeasibility (6) are two such measures.

Step 4 selects the least infeasible $(l,\ e,\ o) = \Pi(\mathcal{W}^r(\mathcal{B}))$ to update $B'^r = B^o \backslash l \cup e$. Each $(a,\ o)$ in $\mathcal{W}^r(\mathcal{B})$ and hence B'^r may result from any $B^o \in \mathcal{B}$. Thus, B'^r may not be adjacent to B^r in $\mathcal{G}$

Three Basis Method

0: <u>INITIALIZE BASES AND BOUNDS ON z^*</u>:
Let $R = \{p,\ d,\ i\}$ be the allowed basis types. (In the two basis method, $R = \{p,\ d\}$.)
Let $\emptyset \neq \mathcal{B} \subseteq \cup_{r \in R} B^r$ be the set of initial bases; e.g., $\mathcal{B} = \{B^p\},\ \{B^d\}$, or $\{B^p,\ B^d\}$.
Initialize lower bound $\underline{z}$ and upper bound $\overline{z}$ on the optimal objective value z^*:

$$\underline{z} = \max\{-\infty,\ \max_{B \in \mathcal{B}:\ L(B)=\emptyset} z(B)\} \quad \text{and} \quad \overline{z} = \min\{\infty,\ \min_{B \in \mathcal{B}:\ E(B)=\emptyset} z(B)\}.$$

1: <u>CHECK FOR TERMINATION</u>: If any basis $B \in \mathcal{B}$ is terminal (see (4)), stop. In particular, if $L(B) = \emptyset$ and $E(B) = \emptyset$, terminate with optimal basis $B^* = B$ and optimal objective $z^* = z$.

2: <u>UPDATE BOUNDS ON z^*</u>: From each $B \in \mathcal{B}$, evaluate the primal (resp. dual) feasibility of all pivots $\mathcal{C}(B)$ so as to maximize $\underline{z}$ (resp. minimize $\overline{z}$) and close duality gap $\overline{z} - \underline{z}$ from both ends:

$$\underline{z} = \max\left\{\underline{z},\ \max_{B \in \mathcal{B},\ a \in \mathcal{C}(B):\ L_a(B)=\emptyset} z_a(B)\right\} \quad \text{and} \quad \overline{z} = \min\left\{\overline{z},\ \min_{B \in \mathcal{B},\ a \in \mathcal{C}(B):\ E_a(B)=\emptyset} z_a(B)\right\}.$$

3: <u>FILTER AND CLASSIFY PIVOTS</u>:
Let $O = \{o \in R :\ B^o \in \mathcal{B}\}$ be the set of originating basis types in $\mathcal{B}$.
Collect pivots $a \in \mathcal{C}(B^o)$ within bounds $z_a \in [\underline{z},\ \overline{z}]$ along with their originating basis $B^o \in \mathcal{B}$:

$$\mathcal{W}(\mathcal{B}) = \cup_{o \in O} \cup_{a \in \mathcal{C}(B^o)} \{(a,\ o) :\ z_a(B^o) \in [\underline{z},\ \overline{z}]\}.$$

Classify $\mathcal{W}(\mathcal{B})$ into primal feasible, dual feasible and primal-and-dual infeasible pivot sets:

$$\begin{aligned}
\mathcal{W}^p(\mathcal{B}) &= \{(a,\ o) \in \mathcal{W}(\mathcal{B}) :\ L_a(B^o) = \emptyset\} \\
\mathcal{W}^d(\mathcal{B}) &= \{(a,\ o) \in \mathcal{W}(\mathcal{B}) :\ E_a(B^o) = \emptyset\} \\
\mathcal{W}^i(\mathcal{B}) &= \{(a,\ o) \in \mathcal{W}(\mathcal{B}) :\ L_a(B^o) \neq \emptyset \text{ and } E_a(B^o) \neq \emptyset\}.
\end{aligned}$$

4: <u>SELECT PIVOTS AND UPDATE BASES</u>:
Initialize the set of new bases $\mathcal{B}' = \emptyset$.
for $r \in R$
  if $\mathcal{W}^r(\mathcal{B}) \neq \emptyset$, i.e., if there are any candidate pivots for a new basis of type r, **then**
    Select least infeasible pivot-and-originating basis $(l,\ e,\ o) = \Pi(\mathcal{W}^r(\mathcal{B}))$.
    Set new basis $B'^r = B^o \backslash l \cup e$ and $\mathcal{B}' = \mathcal{B}' \cup \{B'^r\}$.

(Among all primal feasible bases with non-decreasing $z = \underline{z}$ adjacent to any $B^o \in \mathcal{B}$, this selects the least infeasible B'^p to update B^p. Likewise, it updates B^d to the least infeasible B'^d with non-increasing $z = \overline{z}$, and B^i to the least infeasible B'^i with non-monotone $z \in [\underline{z},\ \overline{z}]$.)

Let $\mathcal{B} = \mathcal{B}'$. Update the dictionary $(b,\ D,\ z,\ c)$ in (2) for each $B \in \mathcal{B}$ and **goto** step 1.

Least infeasible pivot selection rule $\Pi(\underline{\mathcal{W}})$

For $(a,\ o) \in \underline{\mathcal{W}}$, i.e., for pivot $a \in \mathcal{C}(B^o)$, we may compute infeasibility π_{ao} as follows:

$$\pi_{ao} = -\sum_{p \in L_a(B^o)} b_p^a + \sum_{p \in E_a(B^o)} c_p^a \qquad \text{or} \qquad \pi_{ao} = \max\left(\max_{p \in L_a(B^o)} -b_p^a,\ \max_{p \in E_a(B^o)} c_p^a\right) \tag{6}$$

The least infeasibility with respect to pivot-and-basis set $\underline{\mathcal{W}}$ is $\pi' = \min_{(a,\ o) \in \underline{\mathcal{W}}} \pi_{ao}$.
return $\operatorname{argmin}_{(a,\ o) \in \underline{\mathcal{W}}} \pi_{ao}$, i.e., return the least infeasible $(a,\ o) \in \underline{\mathcal{W}} :\ \pi_{ao} = \pi'$.

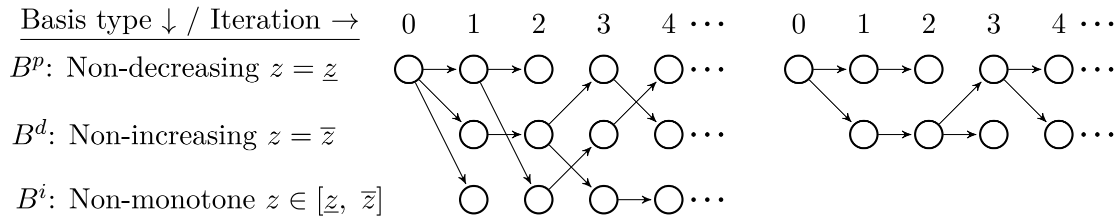

Figure 3: Three basis and two basis pivot paths

denoting a criss-cross path visiting any basis type leading up to B'^r. Any current $B^o \in \mathcal{B}$ may be cut off, or may yield one or more new $B'^r \in \mathcal{B}'$ resulting in interconnected pivot paths; see figure 3.

Objective $z = \underline{z}$ is non-decreasing in B^p, $z = \overline{z}$ is non-increasing in B^d while $z \in [\underline{z},\ \overline{z}]$ may be non-monotone in B^i. Infeasibility $\pi(B^r)$ may be non-monotone along the path of each $B^r,\ r \in R$.

Pivot rule $\Pi(\underline{\mathcal{W}})$ selects the primal (resp. dual) feasible pivot in $\mathcal{W}^p(\mathcal{B})$ (resp. $\mathcal{W}^d(\mathcal{B})$) with the least dual (resp. primal) infeasibility. Absent any primal (resp. dual) inbound pivot, it selects the least infeasible primal (resp. dual) simplex pivot that maximizes $\underline{z}$ (resp. minimizes $\overline{z}$). This is the largest improvement simplex rule for $R = \{p\}$ (resp. $R = \{d\}$); its complexity is exponential [19].

3.2 Two Basis Method

For $R = \{p,\ d\}$, a two basis variant maintains B^p, B^d, or both per iteration to traverse interlinked, non-adjacent, criss-cross paths. It ignores B^i and $\mathcal{W}^i$, losing primal and/or dual inbound pivots from B^i. To improve $\underline{z}$ (resp. $\overline{z}$) and find B'^p (resp. B'^d), it only considers monotone primal (resp. dual) simplex pivots from B^p (resp. B^d) and primal (resp. dual) inbound pivots from B^d (resp. B^p).

3.3 One Basis Methods

Restricting both methods to work with one customary basis $\mathcal{B} = \{B\}$ yields variants that explore its neighbors in $\mathcal{G}$ to try to improve $\underline{z}$ and $\overline{z}$ with non-monotone $z \in [\underline{z},\ \overline{z}]$. Initialized with $B = B^p$ or B^d, one selects $(l,\ e,\ o) = \Pi(\mathcal{W}(B))$ in step 4 to find $B' = B \backslash l \cup e$. The other feasible pivot method selects $(l,\ e,\ o) = \Pi(\mathcal{W}^p(B) \cup \mathcal{W}^d(B))$. The former (resp. latter) walks a criss-cross path visiting intermingled least infeasible bases B'^p, B'^d or B'^i (resp. B'^p or B'^d) in any order.

Breaking Ties Pivot set $\underline{\mathcal{W}}$ in $\Pi(\underline{\mathcal{W}})$ may contain multiple tied least infeasible pivots $\mathcal{W}' = \{(a,\ o) \in \underline{\mathcal{W}} : \pi_{ao} = \pi'\}$. To differentiate within $\mathcal{W}'$ in $\Pi(\mathcal{W}^r(\mathcal{B}))$, $\Pi(\mathcal{W}(B))$ and $\Pi(\mathcal{W}^p(B) \cup \mathcal{W}^d(B))$ of the above methods, we refine $\Pi(\underline{\mathcal{W}})$ to a multi-basis least index rule $\text{LIR}(\Pi(\underline{\mathcal{W}}))$:

Least index rule LIR$(\mathcal{W}')$

Let $h = \min(\{e :\ D_{le}(B^o) < 0,\ (l,\ e,\ o) \in \mathcal{W}'\} \cup \{l :\ D_{le}(B^o) > 0,\ (l,\ e,\ o) \in \mathcal{W}'\})$.
Let $g = \min(\{l :\ D_{lh}(B^o) < 0,\ (l,\ h,\ o) \in \mathcal{W}'\} \cup \{e :\ D_{he}(B^o) > 0,\ (h,\ e,\ o) \in \mathcal{W}'\})$.
if $\exists$ primal pivot $(g,\ h,\ o) \in \mathcal{W}'$ with $D_{gh}(B^o) < 0$ for any $o \in O$ **then**
  return $(g,\ h,\ o)$.
return dual pivot $(h,\ g,\ o) \in \mathcal{W}'$ with $D_{hg}(B^o) > 0$ for any $o \in O$.

Here, h is the least index in $E(B^o)$ (or resp. $L(B^o)$) with a primal (resp. dual) pivot in $\mathcal{W}'$. A fixed preference order, say, $p \succ d \succ i$, breaks ties in O. $\text{LIR}(\mathcal{C}(B))$ is the finite combinatorial least index rule [44] which is exponential [33]. Using $\text{LIR}(\Pi(\underline{\mathcal{W}}))$, it appears the three basis method and its one basis variant may be initialized with any basis. The least infeasibility and least index tie breakers help curb cycling if there is no improvement in $\underline{z}$ of B^p, $\overline{z}$ of B^d, and $[\underline{z},\ \overline{z}]$ of B^i and B.

4 Feasibility Tests

The feasibility tests below check the primal (resp. dual) feasibility of multiple pivots by column (resp. row): explicitly computing L_a (resp. E_a) for each $a \in \mathcal{C}$ is costlier. From (5a) and (5e), let

$$\dot{b}_{ij} = b^{ij}_j = -b_i/D_{ij} \quad \text{and} \quad \dot{c}_{ij} = -c^{ij}_i = -c_j/D_{ij}, \quad \forall i \in B,\ j \in N,\ D_{ij} \neq 0.$$

For pivot $a = (i,\ j)$ to be primal (resp. dual) feasible, from (5b) (resp. (5f)), we need to check if $b^a_p = b_p + b^a_j D_{pj} \geq 0,\ \forall p \in B \backslash i$ (resp. $c^a_q = c_q - c^a_i D_{iq} \leq 0,\ \forall q \in N \backslash j$) where $b^a_j \geq 0$ (resp. $c^a_i \leq 0$):

b_p	$+\, b^a_j$	D_{pj}	$= b^a_p$
$\oplus$	$\oplus$	$\oplus$	$\oplus$
$\oplus$	$\bigcirc$	$-$	$\oplus$
$+$	$+$	$-$	?
$-$	$+$	$+$	?
$-$	$\oplus$	$\ominus$	$-$
$\bigcirc$	$+$	$-$	$-$
$-$	$\bigcirc$	$+$	$-$

$+$:	> 0
$\oplus$:	≥ 0
$\bigcirc$:	$= 0$
$\ominus$:	≤ 0
$-$:	< 0
?:	$\lesseqgtr 0$

c_q	$-\, c^a_i$	D_{iq}	$= c^a_q$
$\ominus$	$\ominus$	$\ominus$	$\ominus$
$\ominus$	$\bigcirc$	$+$	$\ominus$
$-$	$-$	$+$	?
$+$	$-$	$-$	?
$+$	$\ominus$	$\oplus$	$+$
$\bigcirc$	$-$	$+$	$+$
$+$	$\bigcirc$	$-$	$+$

Define the following for $j \in N$ and $i \in B$. Further, let $R_E = \cup_{j \in E} R_j$ and $C_L = \cup_{i \in L} C_i$.

$$R_j = \{i \in B:\ D_{ij} < 0,\ z_{ij} \in [\underline{z},\ \overline{z}]\} \qquad C_i = \{j \in N:\ D_{ij} > 0,\ z_{ij} \in [\underline{z},\ \overline{z}]\}$$

$$I^{<>}_j = \{i \in B:\ b_i < 0,\ D_{ij} > 0\} \qquad J^{><}_i = \{j \in N:\ c_j > 0,\ D_{ij} < 0\}$$

$$I^{<=}_j = \{i \in B:\ b_i < 0,\ D_{ij} = 0\} \qquad J^{>=}_i = \{j \in N:\ c_j > 0,\ D_{ij} = 0\}$$

$$I^{<}_j = \{i \in B:\ D_{ij} < 0\} \qquad J^{>}_i = \{j \in N:\ D_{ij} > 0\}$$

$$\overline{b}_j = \max\{0,\ \max_{i \in I^{<>}_j} \dot{b}_{ij}\} \qquad \overline{c}_i = \max\{0,\ \max_{j \in J^{><}_i} \dot{c}_{ij}\}$$

$$\underline{b}_j = \min\{\infty,\ \min_{i \in I^{<}_j} \dot{b}_{ij}\}. \qquad \underline{c}_i = \min\{\infty,\ \min_{j \in J^{>}_i} \dot{c}_{ij}\}.$$

For $j \in E$ (resp. $i \in L$), the *primal (resp. dual) feasibility test* tries to find primal inbound primal (resp. dual inbound dual) pivots if $L \neq \emptyset$ (resp. $E \neq \emptyset$); else it simplifies to the primal (resp. dual) simplex minimum ratio test to find primal (resp. dual) simplex pivots. For $j \in C_L$ (resp. $i \in R_E$), it looks for primal inbound dual (resp. dual inbound primal) pivots; see figure 1.

After evaluating all feasibility tests, we remove primal or dual feasible pivots from $\mathcal{I}$ to compute $\mathcal{W}^i$: $\mathcal{W}^i = \mathcal{I} \backslash (\mathcal{W}^p \cup \mathcal{W}^d)$. A final filter updates $\mathcal{W}^r = \{(a,\ o) \in \mathcal{W}^r:\ z_a(B^o) \in [\underline{z},\ \overline{z}]\},\ \forall r \in R$.

Initialize infeasibility $\pi_a = 0,\ \forall a \in \mathcal{C}(B)$. Let $\beta^i_{pj} = (\dot{b}_{pj} - \dot{b}_{ij}) D_{pj}$ and $\kappa^j_{iq} = (\dot{c}_{ij} - \dot{c}_{iq}) D_{iq}$. We may update π_{ij} using (7a) in primal feasibility tests for $j \in E$ and $i \in R_j$ and for $j \in C_L$ and $i \in I^{<>}_j$, and using (7b) in dual feasibility tests for $i \in L$ and $j \in C_i$ and for $i \in R_E$ and $j \in J^{><}_i$.

$$\pi_{ij} = \pi_{ij} + \sum_{p \in I^{<>}_j:\ \dot{b}_{pj} > \dot{b}_{ij}} \beta^i_{pj} + \sum_{p \in I^{<}_j:\ \dot{b}_{ij} > \dot{b}_{pj}} \beta^i_{pj} - \sum_{p \in I^{<=}_j} b_p + \max\{0,\ -\dot{b}_{ij}\} \tag{7a}$$

$$\pi_{ij} = \pi_{ij} + \sum_{q \in J^{><}_i:\ \dot{c}_{iq} > \dot{c}_{ij}} \kappa^j_{iq} + \sum_{q \in J^{>}_i:\ \dot{c}_{ij} > \dot{c}_{iq}} \kappa^j_{iq} + \sum_{q \in J^{>=}_i} c_q + \max\{0,\ -\dot{c}_{ij}\}. \tag{7b}$$

In [6], a pivot rule uses dual inbound primal (resp. primal inbound dual) pivot tests to alternate from B^p (resp. B^d) to an adjacent B'^d (resp. B'^p) with the minimum $\overline{z}$ (resp. maximum $\underline{z}$).

Open Questions Are the proposed pivot methods finite? If so, what is their complexity? For the proposed ideas to be practically viable, it is necessary to quickly find simplex or inbound

if $j \in E$ and $I_j^{:<} = \emptyset$ **then**
The dual LP is inconsistent. Stop.
if $I_j^{<=} = \emptyset$ and $\bar{b}_j \le \underline{b}_j$ **then**
if $j \in E$ **then** ▷ Check primal pivots
$\underline{z} = \max(\underline{z},\ z + \underline{b}_j c_j)$
$\mathcal{W}^p = \mathcal{W}^p \cup \bigcup_{i \in R_j:\ \dot{b}_{ij} = \underline{b}_j} (i,\ j)$
$\mathcal{I} = \mathcal{I} \cup \bigcup_{i \in R_j:\ \dot{b}_{ij} < \bar{b}_j \text{ or } \underline{b}_j < \dot{b}_{ij}} (i,\ j).$
if $j \in C_L$ **then** ▷ Check dual pivots
if $I_j^{<>} \neq \emptyset$ **then** $\underline{z} = \max(\underline{z},\ z + \bar{b}_j c_j)$
$\mathcal{W}^p = \mathcal{W}^p \cup \bigcup_{i \in I_j^{<>}:\ \dot{b}_{ij} = \bar{b}_j} (i,\ j)$
$\mathcal{I} = \mathcal{I} \cup \bigcup_{i \in I_j^{<>}:\ \dot{b}_{ij} < \bar{b}_j \text{ or } \underline{b}_j < \dot{b}_{ij}} (i,\ j).$
else
if $j \in E$ **then** $\mathcal{I} = \mathcal{I} \cup \bigcup_{i \in R_j} (i,\ j)$
if $j \in C_L$ **then** $\mathcal{I} = \mathcal{I} \cup \bigcup_{i \in I_j^{<>}} (i,\ j)$

Primal feasibility test $j \in E \cup C_L$

if $i \in L$ and $J_i^{:>} = \emptyset$ **then**
The primal LP is inconsistent. Stop.
if $J_i^{>=} = \emptyset$ and $\bar{c}_i \le \underline{c}_i$ **then**
if $i \in L$ **then** ▷ Check dual pivots
$\bar{z} = \min(\bar{z},\ z + b_i \underline{c}_i)$
$\mathcal{W}^d = \mathcal{W}^d \cup \bigcup_{j \in C_i:\ \dot{c}_{ij} = \underline{c}_i} (i,\ j)$
$\mathcal{I} = \mathcal{I} \cup \bigcup_{j \in C_i:\ \dot{c}_{ij} < \bar{c}_i \text{ or } \underline{c}_i < \dot{c}_{ij}} (i,\ j).$
if $i \in R_E$ **then** ▷ Check primal pivots
if $J_i^{><} \neq \emptyset$ **then** $\bar{z} = \min(\bar{z},\ z + b_i \bar{c}_i)$
$\mathcal{W}^d = \mathcal{W}^d \cup \bigcup_{j \in J_i^{><}:\ \dot{c}_{ij} = \bar{c}_i} (i,\ j)$
$\mathcal{I} = \mathcal{I} \cup \bigcup_{j \in J_i^{><}:\ \dot{c}_{ij} < \bar{c}_i \text{ or } \underline{c}_i < \dot{c}_{ij}} (i,\ j).$
else
if $i \in L$ **then** $\mathcal{I} = \mathcal{I} \cup \bigcup_{j \in C_i} (i,\ j)$
if $i \in R_E$ **then** $\mathcal{I} = \mathcal{I} \cup \bigcup_{j \in J_i^{><}} (i,\ j)$

Dual feasibility test $i \in L \cup R_E$

pivots that improve $\underline{z}$ or $\bar{z}$, even if not maximally, balancing neighborhood diversity and time per iteration. A fast iteration may test one $j \in E \cup C_L$ (resp. $i \in L \cup R_E$) per B to preserve or attain primal (resp. dual) feasibility and use a simpler selection criterion. Despite scope for parallelization, its sparse implementation would be costlier than an efficient revised simplex iteration even for LPs with $0 \le x \le \infty$. It remains to be seen if this is offset by a better convergence rate and runtime.

Acknowledgments The author would like to thank Professors Bob Bulfin, Komei Fukuda, Irv Lustig, Nimrod Megiddo, Katta G. Murty, Michael Saunders, Tamás Terlaky, Craig Tovey, Robert Vanderbei, Shuzhong Zhang and Günter Ziegler for their kind response to his request (in 2000) for comments and assistance to help find a cycling or exponential example for the one basis methods.

References

[1] N. Amenta and G.M. Ziegler. Deformed products and maximal shadows. *Advances in Discrete and Computational Geometry, Contemporary Mathematics*, 223:57–90, 1998.

[2] M. Bárász and S. Vempala. A new approach to strongly polynomial linear programming. In *Proc. Innov. Comp. Sci.*, pages 42–48, 2010.

[3] R.E. Bixby. Solving real-world linear programs: A decade and more of progress. *Oper. Res.*, 50:3–15, 2002.

[4] K.H. Borgwardt. The average number of pivot steps required by the simplex method is polynomial. *Zeitschrift Oper. Res.*, 26:157–177, 1982.

[5] D. Cardoso and J. Clímaco. The generalized simplex method. *OR Letters*, 12:337–348, 1992.

[6] H.-D. Chen, P.M. Pardalos, and M.A. Saunders. The simplex algorithm with a new primal and dual pivot rule. *OR Letters*, 16:121–127, 1994.

[7] S. Chubanov. A strongly polynomial algorithm for linear systems having a binary solution. *Math. Prog.*, 134:533–570, 2012.

[8] V. Chvátal. *Linear Programming.* W.H. Freeman, New York, 1983.

[9] L. Cooper and J. Kennington. Nonextreme point solution strategies for linear programs. *Naval Res. Log. Qtrly.*, 26:447–461, 1979.

[10] G.B. Dantzig. Maximization of a linear function of variables subject to linear inequalities. In *T.C. Koopmans, ed., Activity Analysis of Production and Allocation, 339-347.* Wiley, 1951.

[11] G.B. Dantzig. *Linear Programming and Extensions.* Princeton University Press, 1963.

[12] K. Dosios and K. Paparrizos. Resolution of the problem of degeneracy in a primal and dual simplex algorithm. *OR Letters*, 20:45–50, 1997.

[13] O. Friedmann, T.D. Hansen, and U. Zwick. Subexponential lower bounds for randomized pivoting rules for the simplex algorithm. In *Proc. 43rd ACM STOC*, pages 283–292, 2011.

[14] K. Fukuda and T. Terlaky. Criss-cross methods: A fresh view on pivot algorithms. *Math. Prog.*, 79:369–395, 1997.

[15] K. Fukuda and T. Terlaky. On the existence of a short admissible pivot sequence for feasibility and linear optimization problems. *Pure Math. Appl.*, 10:431–447, 1999.

[16] D.H. Gale, H.W. Kuhn, and A.W. Tucker. Linear programming and the theory of games. In *T.C. Koopmans, ed., Activity Analysis of Production and Allocation, 317-329.* Wiley, 1951.

[17] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization.* Springer, 1988.

[18] E. Hazan and N. Megiddo. The "arrangement method" for linear programming is equivalent to the phase one method. Research report, IBM, 2007.

[19] R.G. Jeroslow. The simplex algorithm with the pivot rule of maximizing criterion improvement. *Disc. Math.*, 4:367–377, 1973.

[20] G. Kalai. A sub-exponential randomized simplex algorithm. In *Proc. 24th ACM STOC*, pages 475–482, 1992.

[21] G. Kalai and D.J. Kleitman. A quasi-polynomial bound for the diameter of graphs of polyhedra. *Bull. Amer. Math. Soc.*, 26:315–316, 1992.

[22] N.K. Karmarkar. A new polynomial-time algorithm for linear programming. *Combinatorica*, 4:373–395, 1984.

[23] J.A. Kelner and D.A. Spielman. A randomized polynomial-time simplex algorithm for linear programming. In *Proc. 38th ACM STOC*, pages 51–60, 2006.

[24] L.G. Khachian. A polynomial algorithm for linear programming. *Soviet Math. Dokl.*, 20:191–194, 1979.

[25] V. Klee and G.J. Minty. How good is the simplex algorithm? In O. Shisha, editor, *Inequalities*, pages 159–175. Academic Press, 1972.

[26] V. Koltun. The arrangement method for linear programming. Preprint, Stanford Univ., 2006.

[27] J. De Loera. New insights into the complexity and geometry of linear optimization. *Optima*, 87:1–13, 2011.

[28] J. Matoušek, M. Sharir, and E. Welzl. A sub-exponential bound for linear programming. In *Proc. 8th ACM Symp. on Comp. Geometry*, pages 1–8, 1992.

[29] N. Megiddo. Towards a genuinely polynomial algorithm for linear programming. *SIAM J. Comp.*, 12:347–353, 1983.

[30] N. Megiddo. Linear programming in linear time when the dimension is fixed. *JACM*, 31:114–127, 1984.

[31] K.G. Murty. Computational complexity of parametric linear programming. *Math. Prog.*, 19:213–219, 1980.

[32] J. von Neumann. Discussion of a maximum problem. Working paper, Princeton, 1947.

[33] C. Roos. An exponential example for Terlaky's pivoting rule for the criss-cross simplex method. *Math. Prog.*, 46:78–94, 1990.

[34] F. Santos. A counterexample to the Hirsch conjecture. *Ann. Math.*, 176:383–412, 2012.

[35] A. Schrijver. *Theory of Linear and Integer Programming.* Wiley, New York, 1986.

[36] H.D. Sherali, A.L. Soyster, and S.G. Baines. Nonadjacent extreme point methods for solving linear programs. *Naval Res. Log. Qtrly.*, 30:145–161, 1983.

[37] S. Smale. On the average number of steps of the simplex method of linear programming. *Math. Prog.*, 27:241–262, 1983.

[38] S. Smale. Mathematical problems for the next century. *Math. Intelligencer*, 20:7–15, 1998.

[39] D. Spielman and S.-H. Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *JACM*, 51:385–463, 2004.

[40] R.P. Stanley. *An Introduction to Hyperplane Arrangements.* AMS, Providence, RI, 2007.

[41] J.E. Strum. *Introduction to Linear Programming.* Holden-Day, San Francisco, 1972.

[42] E. Tardos. A strongly polynomial minimum cost circulation algorithm. *Combinatorica*, 5:247–255, 1985.

[43] E. Tardos. A strongly polynomial algorithm to solve combinatorial linear programs. *Oper. Res.*, 34:250–256, 1986.

[44] T. Terlaky. A convergent criss-cross method. *Optimization*, 16:683–690, 1985.

[45] T. Terlaky and S. Zhang. Pivot rules for linear programming: A survey on recent theoretical developments. *Annals Oper. Res.*, 46:203–233, 1993.

[46] J.F. Traub and H. Wozniakowski. Complexity of linear programming. *OR Letters*, 1:59–62, 1982.

[47] Y. Ye. The simplex and policy-iteration methods are strongly polynomial for the Markov decision problem with a fixed discount rate. *Math. Oper. Res.*, 36:593–603, 2011.

[48] S. Zionts. The criss-cross method for solving linear programming problems. *Mgmt. Sci.*, 15:426–445, 1969.