

Vector Space Decomposition for Solving Large-Scale Linear Programs

Jean Bertrand Gauthier,^a Jacques Desrosiers,^a Marco E. Lübbecke^b

^aHEC Montréal and GERAD, Montréal, Canada H3T 2A7; ^bLehrstuhl für Operations Research, RWTH Aachen University, D-52072 Aachen, Germany

Contact: jean-bertrand.gauthier@hec.ca (JBG); jacques.desrosiers@hec.ca,  <http://orcid.org/0000-0003-0719-1500> (JD); marco.luebbecke@rwth-aachen.de,  <http://orcid.org/0000-0002-2635-0522> (MEL)

Received: September 8, 2016

Revised: March 20, 2017; September 28, 2017

Accepted: December 23, 2017

Published Online in *Articles in Advance*: August 2, 2018

Subject Classifications: linear programming; algorithms

Area of Review: Optimization

<https://doi.org/10.1287/opre.2018.1728>

Copyright: © 2018 INFORMS

Abstract. We develop an algorithmic framework for linear programming guided by dual optimality considerations. The solution process moves from one feasible solution to the next according to an exchange mechanism that is defined by a direction and a resulting step size. Part of the direction is obtained via a pricing problem devised in primal and dual forms. From the dual perspective, one maximizes the minimum reduced cost that can be achieved from splitting the set of dual variables in two subsets: one being fixed while the other is optimized. From the primal perspective, this amounts to selecting a nonnegative combination of variables entering the basis. The direction is uniquely complemented by identifying the affected basic variables, if any.

The framework is presented in a generic format motivated by and alluding to concepts from network flow problems. It specializes to a variety of algorithms, several of which are well known. The most prominent is the primal simplex algorithm where all dual variables are fixed: this results in the choice of a single entering variable commonly leading to degenerate pivots. At the other extreme, we find an algorithm for which all dual variables are optimized at every iteration. Somewhere in between these two extremes lies the improved primal simplex algorithm for which one fixes the dual variables associated with the nondegenerate basic variables and optimizes the remaining ones. The two last variants both bestow a pricing problem providing necessary and sufficient optimality conditions. As a result, directions yielding strictly positive step sizes at every iteration are also issued from these pricing steps. These directions move on the edges of the polyhedron for the latter while the former can also identify interior directions.

Funding: Jacques Desrosiers acknowledges the National Science and Engineering Research Council of Canada [Grant RGPIN/04401-2014] and the HEC Montréal Foundation for their financial support. Jean Bertrand Gauthier acknowledges the GERAD—Group for Research in Decision Analysis (of Montréal, Québec, Canada)—for its financial support.

Keywords: primal simplex algorithm • column generation • degeneracy • residual problem • optimized reduced costs • cycles • positive step size algorithms • vector space

1. Introduction

Degeneracy is a critical performance issue when solving linear programs with the simplex method in practice.

Dantzig and Thapa (2003, p. 167) suggest “that pivot-selection criteria should be designed to seek feasible solutions in directions away from degenerate and ‘near’-degenerate basic feasible solutions, or better yet, driven by dual feasibility considerations.” We revisit that statement with a different interpretation: when trying to avoid primal infeasible directions, one should consider pivot-selection (or pricing) rules that are also guided by *dual optimality*.

While Dantzig’s pivot rule accurately measures the improvement rate of the objective function, the influence on the affected basic variables is taken for granted for every nonbasic variable unit change. When one realizes that not all affected basic variables *can* actually

move, it becomes clear that the pricing rule suffers from a *visibility problem* in terms of basic variable space. Indeed, null step sizes can only happen when at least one basic variable is valued at one of its bounds. As such, degeneracy is a vicious cycle in the sense that solutions with higher levels of degeneracy are more likely to yield degenerate pivots. Additional computations based on the basis are incorporated in the pricing step of steepest edge and Devex pivot rules (Forrest and Goldfarb 1992, Harris 1973), which helps alleviate, but not eliminate, this shortcoming. We propose the elimination of basic variables leading to degeneracy from consideration in the pricing step: This is sufficient to overcome primal degeneracy. It grants more freedom in the dual variables thus allowing the pricing process flexibility in meeting dual optimality.

We propose a framework, which, given a basic feasible solution, *fixes the values of a subset of dual variables*

and optimizes the remaining ones for maximizing the minimum reduced cost. In the primal interpretation, this results in a pricing problem in which one selects a nonnegative combination of variables entering the basis. The way to divide the dual variables into two subsets relies on the choice of a vector subspace basis in the primal, capitalizing on the actual values taken by the basic variables. This opens a wide spectrum of possibilities. The general scheme resembles a dynamic decomposition like Dantzig and Wolfe's and inspires the name of *vector space decomposition*. It unifies a variety of specialized algorithms for linear and network programs. The most prominent special case is the primal simplex algorithm (PS) where *all* dual variables are fixed: this results in the choice of a single entering variable commonly leading to degenerate pivots. At the other extreme when *no* dual variables are fixed, that is, all dual variables are optimized at every iteration, we find the so-called minimum weighted cycle-canceling algorithm (MWCC) which is strongly polynomial when solving network flow problems (Goldberg and Tarjan 1989). Somewhere in between these two extremes lie the improved primal simplex algorithm (IPS) of Elhallaoui et al. (2011) for which one only fixes the dual variables associated with the nondegenerate basic variables, and the dynamic constraint aggregation method (DCA) of Elhallaoui et al. (2005) specifically designed to overcome degeneracy in the context of set partitioning models when solving the linear relaxation by column generation (see Lübbecke and Desrosiers 2005). Building our unified framework is motivated and enabled by network flow analogies seen in Gauthier et al. (2017). Alluding to Dantzig and Thapa (1997), our proposal can also be interpreted as a very general *dual guided* pivot rule.

Besides having a common theoretical container for such seemingly different algorithms as PS, IPS, DCA, MWCC, and (yet unknown) others, we feel that the framework itself opens a new avenue of generally coping with degeneracy in simplex-like (and related) algorithms. Much more importantly, however, we see an utmost practical perspective for our work. The primal simplex method has strong competitors with the dual simplex and barrier methods (Bixby 2002). In particular in linear programming based branch-and-bound, primal algorithms are not the first choice. In column generation, however, the primal simplex algorithm is inherently encoded in the pricing mechanism, regardless of which algorithm is actually used to reoptimize the linear restricted master programs. Column generation gains more and more significance in solving well-structured very large-scale linear programs from practical applications. As a descendant of the primal simplex method, it inherits all difficulties with degeneracy, and the linear relaxations of combinatorial optimization problems particularly suffer from this. Our

framework offers a general and flexible remedy, and yet, it allows (and benefits from) tailoring to the particular application at hand. Furthermore, it is fortunate that our framework plays particularly well with a number of alternative suggestions to cope with degeneracy or degeneracy-related effects in column generation, such as dual variable stabilization (du Merle et al. 1999, Ben Amor et al. 2009).

The paper is organized as follows. Section 2 takes a close look at the essential components of the framework. Several concepts that partake (or not) in the resolution process of a linear program are examined such as nondegenerate pivots, cycles and directions, and the so-called *residual problem*. Each of these is presented in a separate manner, whereas the last subsection ties everything together. Section 3 builds upon these ties and gives birth to the generic framework. Two propositions are exposed in Section 4, the first determines conditions guaranteeing positive step sizes at every iteration and the second shows that identified directions can be interior rather than along edges. We conclude in Section 5 with our contribution and some research perspectives.

2. Linear Program

Consider the linear program (LP) with lower and upper bounded variables:

$$\begin{aligned} z^* &:= \min c^\top x \\ \text{s.t. } & Ax = b \quad [\pi] \\ & l \leq x \leq u, \end{aligned} \quad (1)$$


where $x, c, l, u \in \mathbb{R}^n$, $b \in \mathbb{R}^m$, $A \in \mathbb{R}^{m \times n}$, and $m < n$. We assume that the matrix A is of full row rank, that LP (1) is feasible, and that z^* is finite. The vector of dual variables $\pi \in \mathbb{R}^m$ associated with the equality constraints appears within brackets on the right-hand side.

2.1. Notation

Vectors and matrices are written in bold face by, respectively, using lower and upper case. We denote by I_r the $r \times r$ identity matrix and by $\mathbf{0}$ (respectively, $\mathbf{1}$) a vector/matrix with all zeros (respectively, ones) entries of contextually appropriate dimension. For an ordered subset $R \subseteq \{1, \dots, m\}$ of row indices and an ordered subset $P \subseteq \{1, \dots, n\}$ of column indices, we denote by A_{RP} the submatrix of A containing the rows and columns indexed by R and P , respectively. We further use standard linear programming notation like $A_B x_B$, the subset of basic columns of A indexed by B multiplied by the vector of basic variables x_B . The index set of nonbasic columns N is used analogously.

In Section 2.2, we formulate the so-called *residual problem*, which allows the construction of an oracle generating feasible directions in Section 2.4. The latter also provides two alternative primal and dual conditions characterizing optimality for linear programs. Finally,

Figure 1. (Color online) Forward and Backward Variables for the Residual Problem



let us embark upon this generic framework in Section 2.5 by analyzing a linear transformation, the goal being to structure the technological constraints.

2.2. Residual Problem

It is a common concept in developing network flow algorithms to use a *residual network* to improve upon some intermediate solution by identifying augmenting flows; see Ahuja et al. (1993). Let us stay in the spirit of network flows and propose analogies for linear programs leading to primal-dual optimality conditions on the so-called residual problem.

We define the *residual problem* $LP(x^k)$ with respect to a given solution x^k at iteration $k \geq 0$ as follows. Each variable x_j , $j \in \{1, \dots, n\}$, in the original LP (1) is replaced by two variables: the forward variable y_j of cost $d_j := c_j$ represents the possible increase $r_j^k := u_j - x_j^k$ of x_j relatively to x_j^k while the backward variable y_{j+n} of cost $d_{j+n} := -c_j$ represents its possible decrease $r_{j+n}^k := x_j^k - l_j$. In fact, the y -variables satisfy $y_j - y_{j+n} = x_j - x_j^k$, $\forall j \in \{1, \dots, n\}$. Moreover, only one from each pair can be used with a positive value, that is, the condition $y_j y_{j+n} = 0$ holds, $\forall j \in \{1, \dots, n\}$; see Figure 1.

Equivalent to LP (1), a formulation for the residual problem $LP(x^k)$ is given as


$$\begin{aligned} z^* &= z^k + \min d^\top y \\ \text{s.t. } & Ky = 0 \quad [\pi] \\ & 0 \leq y \leq r^k, \end{aligned} \quad (2)$$

where $z^k := c^\top x^k$, $d := [d_j]_{j \in \{1, \dots, 2n\}}$ is the cost vector, $y := [y_j]_{j \in \{1, \dots, 2n\}}$ contains the forward and backward variables, their residual capacities given by the upper bounds $r^k := [r_j^k]_{j \in \{1, \dots, 2n\}}$, and the matrix $K := [A, -A] \equiv [k_j]_{j \in \{1, \dots, 2n\}}$ stands to remind us that the kernel (or null space) of this matrix is the set of solutions to $Ky = 0$. A variable fixed to 0 is uninteresting, so $LP(x^k)$ may be written using only the *residual variables*, that is, the y -variables with positive residual capacities within the set $J^k := \{j \in \{1, \dots, 2n\} \mid r_j^k > 0\}$. While this might appear like a trivial statement, it is solely responsible for allowing the identification of directions inducing null step sizes as presented in the next section.

2.3. Directions and Cycles

By neglecting the upper bounds from (2), one obtains a cone whose every extreme ray induces an extreme direction originating from x^k as stated in Definition 1.

Figure 2. (Color online) The Cone $\{y \geq 0 \mid Ky = 0\} \subseteq \mathbb{R}_+^{2n}$ Cut by $w^\top y = 1$



Definition 1 (Bazaraa et al. (1990)). Given an extreme ray $y \in \{Ky = 0, y \geq 0\} \subseteq \mathbb{R}_+^{2n}$, the components of *direction* $\vec{v} \in \mathbb{R}^n$ are computed as differences:

$$\vec{v}_j = y_j - y_{j+n}, \quad \forall j \in \{1, \dots, n\}.$$

To find an improving direction, it then suffices to look within this set of extreme rays. This can be done via an optimization program. However, since optimizing in a cone is inconvenient with respect to any comparative measure (unless the optimal solution is the only extreme point), let us add the normalization constraint $w^\top y = 1$, where $w > 0$ is a vector of arbitrary positive weights. This results in a cut cone where every nonnull extreme point corresponds to an extreme ray. In the same vein, by considering a different weight vector w , the cutting plane would be slanted differently thus producing a modified set of extreme points, yet each of these would remain associated with the same extreme ray. The cone displayed in Figure 2 illustrates these extreme rays.

A *normalized direction* is a direction (abusively speaking by the correspondence between extreme rays and extreme points) for which y satisfies the intersection between the cone and the normalization constraint, that is,

$$y \in \mathcal{N} := \{Ky = 0, y \geq 0, w^\top y = 1\} \subseteq \mathbb{R}_+^{2n}. \quad (3)$$

In the following, Definition 2 introduces the notion of *cycle*, which interchangeably characterizes a normalized direction, whereas Lemma 1 gives us a fundamental property regarding the cost and the reduced cost of such a cycle. Observe that the domain \mathcal{N} does not depend on the current solution x^k . However, by definition of the cone, a cycle may always be followed from x^k with a nonnegative step size $\rho \geq 0$. Definition 3 reconsiders the existence of the residual bounds r^k and characterizes the cycles that feature a strictly positive step size with respect to x^k .

Definition 2. Let $w \in \mathbb{R}^{2n}$ be a vector of strictly positive weights. A *cycle* W is the positive variable support of y in a normalized direction, that is, $W := \{j \in \{1, \dots, 2n\} \mid y_j > 0, y \in \mathcal{N}\}$.

The cost of a cycle W in (2) is computed as $d_W := \sum_{j \in W} d_j y_j$. A *negative cycle* is a cycle with negative cost. The reduced cost of variable $x_j, j \in \{1, \dots, n\}$, is defined as $\bar{c}_j := c_j - \pi^T a_j$ while those of variables y_j and y_{j+n} are, respectively, $\bar{d}_j := c_j - \pi^T a_j = \bar{c}_j$ and $\bar{d}_{j+n} := -c_j - \pi^T(-a_j) = -\bar{c}_j$. The following provides the analogous result from networks that the cost and reduced cost of a cycle are equal.

Lemma 1. *The cost d_W and the reduced cost \bar{d}_W of a cycle W are equal.*

Proof. A cycle W satisfies (3), hence $\sum_{j \in W} k_j y_j = \mathbf{0}$ and $\bar{d}_W = \sum_{j \in W} (d_j - \pi^T k_j) y_j = d_W - \pi^T \sum_{j \in W} k_j y_j = d_W$. \square

Definition 3. An *augmenting cycle* at x^k is a cycle W whose composing y -variables have positive residual capacities, that is, $r_j^k > 0, \forall j \in W$, or equivalently $W \subseteq J^k$.

When in the following we speak of negative cycles, we always refer to negative *augmenting* cycles as per Definition 3 and Lemma 1. Under the previous nomenclature, necessary and sufficient optimality conditions come together in a straightforward manner. In addition to the complementary slackness optimality conditions based on the reduced cost of the original x -variables, we repeat here two alternative conditions characterizing optimality for LP (1).

Proposition 1 (Gauthier et al. (2014, Theorem 4)). *A feasible solution x^k to LP (1) is optimal if and only if the following equivalent conditions are satisfied:*

Complementary slackness: There exists a π such that

$$\begin{aligned} \bar{c}_j > 0 &\Rightarrow x_j^k = l_j; & \bar{c}_j < 0 &\Rightarrow x_j^k = u_j; \\ l_j < x_j^k < u_j &\Rightarrow \bar{c}_j = 0. \end{aligned} \quad (4)$$

Primal: LP(x^k) contains no negative cycle, that is,

$$d_W \geq 0, \quad \text{for every cycle } W \text{ in } LP(x^k). \quad (5)$$

Dual: There exists a π such that the reduced cost of every residual variable of LP(x^k) is nonnegative, that is,

$$\bar{d}_j \geq 0, \quad \forall j \in J^k. \quad (6)$$

2.4. Oracle

To prove the optimality of x^k or improve the current solution, we derive an oracle relying on the identification of cycles. One can be derived from the domain (3) and an objective function that effectively computes the cost (or the reduced cost) of each cycle properly as follows:

$$\min_{y \in \mathcal{N}} d^T y. \quad (7)$$

Since upper bounds are removed from the domain \mathcal{N} , the oracle (7) may unfortunately identify a nonaugmenting cycle if any y -variable with zero residual capacity remains. Keeping track of the residual

variables can be done by partitioning the x -variables according to their current values. To achieve this, let x^k be represented by $(x_{F_i}^k, x_{L_i}^k, x_{U_i}^k)$, where the three subvectors refer to the set of free variables $F := \{j \in \{1, \dots, n\} \mid l_j < x_j^k < u_j\}$, at their lower bounds $L := \{j \in \{1, \dots, n\} \mid x_j^k = l_j\}$, and at their upper bounds $U := \{j \in \{1, \dots, n\} \mid x_j^k = u_j\}$, respectively. Let there be $f := |F|$ such free variables, $0 \leq f \leq n$. Observe that if x^k is basic then $0 \leq f \leq m$. Controlling the presence/absence of the residual variables while solving the oracle can then alternatively be achieved by imposing

$$y_j = 0, \quad \forall j \in U \quad \text{and} \quad y_{j+n} = 0, \quad \forall j \in L. \quad (8)$$

By the primal optimality conditions in Proposition 1, it is possible to improve intermediate solutions using negative cycles until an optimal solution is reached. In this respect, the step size ρ associated with the negative cycle W must satisfy $\rho y_j \leq r_j^k, \forall j \in W$, and this cycle is *canceled* when the step size is equal to $\rho := \min_{j \in W} r_j^k / y_j > 0$.

2.4.1. Primal Simplex Algorithm. Consider a basic solution x^k and the index set of basic variables B within PS. A pivot operation tries to improve the current solution using a nonbasic entering variable, say $x_l, l \in N$. The aftermath of this operation is simplified to a properly selected exiting variable and the associated step size ρ is determined by the ratio test. The ratio test is useful on two counts. It maximally exploits the exchange potential of the entering variable and it maintains a basic solution for x^{k+1} . The mechanic is incredibly simple although it might sometimes render the linear algebra aspect of the pivot nebulous, especially in the context of degeneracy. In this respect, when PS performs a nondegenerate pivot at iteration $k \geq 0$, it goes from vertex x^k represented by a nonoptimal basis to vertex x^{k+1} by moving along an edge (Dantzig and Thapa 2003, Theorem 1.7), a direct consequence of the entering/exiting variable mechanism. In the case of a degenerate pivot, the basis is modified, but the geometrical vertex solution remains the same. In other words, the n -dimensional direction \vec{v} (see Definition 1)

$$\vec{v}_j = \begin{cases} y_j - y_{j+n}, & \forall j \in B \cup \{l\} \\ 0, & \forall j \in N \setminus \{l\} \end{cases} \quad (9)$$

induced by the selected negative reduced cost entering variable x_l leads outside the domain of LP (1) and we do not move. One may want to consider the column a_l of the entering variable as part of the linear span of A_B , defined as $V(A_B) := \{A_B \lambda \mid \lambda \in \mathbb{R}^m\}$. By definition, any m -dimensional column belongs to $V(A_B) = \mathbb{R}^m$ meaning in particular that for any nonbasic entering variable

$$\begin{aligned} \exists! \lambda \in \mathbb{R}^m \text{ such that } \sum_{j \in B} a_j \lambda_j = a_l \\ \text{which works out to } \lambda = A_B^{-1} a_l. \end{aligned} \quad (10)$$

Observe from (9) that a direction in the primal simplex algorithm is not given by the sole entering variable, nor is it limited to the entering/exiting variable couple. It is rather associated with a cycle that combines the entering variable to the affected basic ones. Since the linear combination scalars λ in (10) can take any sign, every column of A_B is implicitly expected to have freedom to move in either direction. This *could* unfortunately turn out to be false when the pivot exercise arrives. This possibility can only arise when a non-basic column is defined by a linear combination containing at least one basic variable at one of its bounds. Indeed, the cycle associated to such an entering variable *might* include a y_j -variable, $j \in B$, with a residual bound of 0, that is, a forward variable $y_j > 0$ with $x_j^k = u_j$ or a backward variable $y_{j+n} > 0$ with $x_j^k = l_j$. The reader may want to compare this with (8) to realize that the degeneracy phenomenon takes an equivalent form in the oracle as well.

2.5. Linear Algebra

Applying the inverse of an arbitrary nonsingular matrix on the equality constraints of LP (1) yields an equivalent system. The goal of the linear transformation T_p^{-1} we propose is to structure the equality constraints, where $P \subseteq B$ is an ordered subset of the indices of the basic variables. The set P induces a *subspace basis* A_p with dimension $p := |P|$. In that case, a subset of p rows within A_p are independent. There then exists a row partition $\begin{bmatrix} A_{RP} \\ A_{SP} \end{bmatrix}$ of A_p such that A_{RP} is a nonsingular matrix of size $p \times p$. For instance, an optimal basic solution to the *restricted* phase I problem

$$\min\{\mathbf{1}^\top \theta \mid A_p x_p + I_m \theta = \mathbf{b}, \theta \geq \mathbf{0}\} \quad (11)$$

identifies a set S of rows by associating $(m - p)$ θ -variables with the used columns of I_m yielding the simplex basis $\begin{bmatrix} A_{RP} & \mathbf{0} \\ A_{SP} & I_{m-p} \end{bmatrix}$, hence the requested row partition of A_p .

Let $V(A_p) := \{A_p \lambda \mid \lambda \in \mathbb{R}^p\}$ be the vector subspace of \mathbb{R}^m spanned by A_p . Because every subset of p linearly independent vectors of $V(A_p)$ can be used as a subspace basis for $V(A_p)$, an alternative set to $\begin{bmatrix} A_{RP} \\ A_{SP} \end{bmatrix}$ is $\begin{bmatrix} I_p \\ M \end{bmatrix}$, where $M := A_{SP} A_{RP}^{-1}$. Together with the $m - p$ independent vectors of $\begin{bmatrix} I_{m-p} \\ \mathbf{0} \end{bmatrix}$, it provides the basis T_p of \mathbb{R}^m and its inverse of particularly simple structure:

$$T_p = \begin{bmatrix} I_p & \mathbf{0} \\ M & I_{m-p} \end{bmatrix} \quad \text{and} \quad T_p^{-1} = \begin{bmatrix} I_p & \mathbf{0} \\ -M & I_{m-p} \end{bmatrix}.$$

Applying T_p^{-1} on the system $Ax = \mathbf{b}$ results in $\bar{A} := T_p^{-1}A$ and $\bar{\mathbf{b}} := T_p^{-1}\mathbf{b}$ as follows:

$$\bar{A} = \begin{bmatrix} A_R \\ A_S - MA_R \end{bmatrix} \quad \text{and} \quad \bar{\mathbf{b}} = \begin{bmatrix} \mathbf{b}_R \\ \mathbf{b}_S - M\mathbf{b}_R \end{bmatrix}. \quad (12)$$

Definition 4 (Gauthier et al. (2016, Proposition 3)). A vector $\mathbf{a} \in \mathbb{R}^m$ (and the associated variable, if any) is *compatible with* A_p if and only if $\bar{\mathbf{a}}_S := \mathbf{a}_S - MA_R = \mathbf{0}$ or, equivalently, $\mathbf{a} \in V(A_p)$. Otherwise, the vector \mathbf{a} is *incompatible*.

Verifying the equivalence is straightforward when decomposing $\begin{bmatrix} a_R \\ a_S \end{bmatrix} = \begin{bmatrix} A_{RP} \\ A_{SP} \end{bmatrix} \lambda$. Checking a column vector for compatibility can therefore be done using methods available from the linear algebra arsenal. Some are more efficient than others depending on the content of the matrix A , probing cases being the network and set partitioning problems that easily permit the verification of the definition; see the *Transformation matrix insight* paragraph at the end of this section. Compatibility can also be determined over the basis A_B in $O(m^2)$ using the *positive edge* rule (see Towhidi et al. 2014), which reduces the matrix multiplication computational penalty with a stochastic argument.

2.5.1. Column Partition. Given a *subspace basis* A_p , let $Q := \{1, \dots, n\} \setminus P$ be an ordered subset of the indices of all the variables outside the subset $P \subseteq B$. This additional column partition is represented by the matrix $A = \begin{bmatrix} A_p & A_Q \end{bmatrix}$. Altogether, we have the row/column partition $A = \begin{bmatrix} A_{RP} & A_{RQ} \\ A_{SP} & A_{SQ} \end{bmatrix}$, where the nonsingular matrix A_{RP} is called the *working basis*. Applying T_p^{-1} on A yields

$$\bar{A} = \begin{bmatrix} I_p & \mathbf{0} \\ -M & I_{m-p} \end{bmatrix} \begin{bmatrix} A_{RP} & A_{RQ} \\ A_{SP} & A_{SQ} \end{bmatrix} = \begin{bmatrix} A_{RP} & A_{RQ} \\ \mathbf{0} & A_{SQ} \end{bmatrix}. \quad (13)$$

If one thinks of T_p as the current primal simplex basis matrix, the ratio test of an entering variable x_l with null entries in $\bar{\mathbf{a}}_{Sl}$ would be performed only on the positive coefficients of $\bar{\mathbf{a}}_{Rl}$ and thus only influence variables related to A_p . This means that all variables associated with A_p and the row set R are *assumed* to be free whereas all variables associated with $\begin{bmatrix} I_{m-p} \\ \mathbf{0} \end{bmatrix}$ and the row set S are *assumed* to be at their bounds. If the set P indeed corresponds to free variables only ($A_p = A_F$), the resulting step size would be positive for sure. In this spirit, the purpose of T_p^{-1} is to induce a partition in \bar{A} to help look for so-called compatible column vectors; see Section 3.

2.5.2. Transformation Matrix Insight. Ultimately, the transformation matrix produces row and column partitions intimately linked together. Depending on the application, the row partition can even be obtained in the midst of selecting the set P by trying to capture the linear dependence of the technological constraints. Network flow and set partitioning problems are such applications; see Figures 3 and 4, respectively.

In network flows, the free arcs forming A_F can be separated in trees forming a forest. The latter is expressed in matrix form in Figure 3(a) and one can then associate a root node to each tree (in bold). Each of these root nodes corresponds to a linear dependent row in A_F thus forming the row partition presented in

Figure 5. Generic Vector Space Decomposition Framework for Linear Programs

- Initialization:** Iteration $k := 0$;
 Feasible solution x^0 ;
 Select the subspace basis A_p ;
- 1: Derive the matrix T_p^{-1} and the row/column partitions $\{R, S\}/\{P, Q\}$ of A ;
 - 2: Divide the vector $y \in \mathbb{R}_+^{2n}$ with the three sets H_p, V_p and Z_p as in (14)–(16);
 - 3: Solve the pricing problem (21) for μ_V^k and the projected cycle W_V^k ;
 - 4: If $\mu_V^k \geq 0$, **terminate** with an optimality certificate for x^k ;
 - 5: Extract the lifted cycle $W_{V \cup H}^k$ from (22)–(23) and the direction \bar{v}^k ;
 - 6: Compute the step size ρ^k (24);
 - 7: Obtain the solution $x^{k+1} := x^k + \rho^k \bar{v}^k$ using the system (25);
 - 8: Update the subspace basis A_p ;
 - 9: Iterate $k := k + 1$;

3.2. Structured Residual Problem

Once an arbitrary subspace basis A_p is selected, the induced matrix T_p^{-1} is derived along with a row/column partition $\{R, S\}/\{P, Q\}$ of matrix A . The only point we shall insist on is the structuring effect of T_p^{-1} . The same can obviously be observed in the residual problem $LP(x^k)$. Let us divide the vector $y \in \mathbb{R}_+^{2n}$ according to P with the three sets

$$H_p \equiv H_p(x^k) := \bigcup_{j \in P} \{j, j+n\} \quad (14)$$

$$V_p \equiv V_p(x^k) := J^k \setminus H_p \quad (15)$$

$$Z_p \equiv Z_p(x^k) := \{1, \dots, 2n\} \setminus (H_p \cup V_p), \quad (16)$$

where the y -variables in the set H_p are *hidden* from the oracle, the residual variables in V_p are rather *visible* in the oracle, whereas the remaining ones (with null residual upper bounds) in Z_p remain at zero. The residual y -variables are obviously exhaustively considered within H_p and V_p . However, observe that while V_p contains only residual variables, H_p may or may not. That is, if there exists a $j \in P$ such that either $j \in L$ or $j \in U$, then y_{j+n} or, respectively, y_j has a null residual capacity. Discarding the variables in Z_p from the residual problem, that is, $y_j^k = 0, \forall j \in Z_p$, the formulation (2) can then be rewritten as

$$\begin{aligned} z^* &= z^k + \min \sum_{j \in H_p} d_j y_j + \sum_{j \in V_p} d_j y_j \\ \text{s.t.} \quad & \sum_{j \in H_p} k_{Rj} y_j + \sum_{j \in V_p} k_{Rj} y_j = 0 \quad [\psi_R] \\ & \sum_{j \in V_p} \bar{k}_{Sj} y_j = 0 \quad [\psi_S] \\ & 0 \leq y_j \leq r_j^k, \quad \forall j \in H_p \cup V_p, \end{aligned} \quad (17)$$

where k_j is the j th column vector of K , $\bar{K} := T_p^{-1}K$, and $\psi^T = [\psi_R^T, \psi_S^T]$ is the vector of dual variables of the transformed system. The original dual vector π can be retrieved from ψ using the expression $\pi^T = \psi^T T_p^{-1}$, that is, $[\pi_R^T, \pi_S^T] = [\psi_R^T - \psi_S^T M, \psi_S^T]$.

3.3. Pricing Oracle: Cycle and Direction

The pricing problem exploits the structure in (17) and derives an oracle based on the resulting transformation. The oracle is presented in both primal and dual forms, each having its own interpretation. Let us start with the dual form that is derived by trying to meet the necessary and sufficient optimality conditions. That is, if the assumed free variables in the set P are at an optimal value, the dual variables of π , or those of ψ , must impose a reduced cost of zero on these variables [complementary slackness conditions (4)]:

$$\begin{aligned} 0 &= \bar{c}_P^T = c_P^T - \psi_R^T A_{RP} \\ &= c_P^T - (\pi_R^T + \pi_S^T M) A_{RP} = c_P^T - \pi_R^T A_{RP} - \pi_S^T A_{SP}. \end{aligned} \quad (18)$$

Furthermore, if the current solution x^k is optimal, there exists a dual vector ψ_S such that the smallest reduced cost of the remaining variables in V_p , say μ_V , is nonnegative [dual condition (6)]. This verification can be done with the linear program

$$\begin{aligned} \max \quad & \mu_V \\ \text{s.t.} \quad & \mu_V \leq \bar{d}_j = d_j - \psi_R^T k_{Rj} - \psi_S^T \bar{k}_{Sj} \quad [y_j] \quad \forall j \in V_p, \end{aligned} \quad (19)$$

where the vector $\psi_R^T = c_P^T A_{RP}^{-1}$ is fixed by (18), whereas the vector ψ_S^T is part of the optimization so as to maximize the minimum reduced cost. More generally, one can use the scalars $w_j > 0, \forall j \in V_p$, and maximize the minimum normalized reduced cost (see Section 3.3.1):

$$\begin{aligned} \max \quad & \mu_V \\ \text{s.t.} \quad & \mu_V \leq \frac{\bar{d}_j}{w_j} = \frac{1}{w_j} (d_j - \psi_R^T k_{Rj} - \psi_S^T \bar{k}_{Sj}) \\ & [y_j] \quad \forall j \in V_p. \end{aligned} \quad (20)$$

Dualizing (20), we obtain the primal form that comprises $m - p + 1$ constraints:

$$\begin{aligned} \min \quad & \sum_{j \in V_p} \bar{d}_j y_j \\ \text{s.t.} \quad & \sum_{j \in V_p} \bar{k}_{Sj} y_j = 0 \quad [\psi_S] \\ & \sum_{j \in V_p} w_j y_j = 1 \quad [\mu_V] \\ & y_j \geq 0, \quad \forall j \in V_p, \end{aligned} \quad (21)$$

where $\bar{d}_j := d_j - \psi_R^T k_{Rj} = d_j - c_P^T A_{RP}^{-1} k_{Rj}, \forall j \in V_p$. The oracle interpretation is done through the primal-dual pair (20)–(21). It brings together the negative cycles and T_p^{-1} .

3.3.1. Oracle Interpretation. First of all, the formulation (21) is always feasible unless x^0 is the only feasible solution of LP (1). Reciprocally, (20) is always feasible although unbounded in the exception case. Note that we can ensure that the primal-dual pricing system is feasible/bounded if we write the normalization

constraint as a less-than-or-equal inequality or equivalently one imposes $\mu_V \leq 0$. Furthermore, weight values of w used for the normalization constraint can be set in stone or updated dynamically. In the former case, think of the all-ones vector typically used in network flows (see Gauthier et al. 2015) or the norm based weights such as $w_j = w_{j+n} = \|a_j\|$, $\forall j \in \{1, \dots, n\}$, which makes the ratio \bar{d}_j/w_j impervious to the scaling of variable x_j . In the latter case, dynamic weight choices can also be made to help steer the pricing problem toward or away from certain solutions; see Rosat et al. (2017a, b) for several alternatives that have been particularly successful for solving set partitioning problems using the *integral simplex using decomposition* algorithm (ISUD) (Zaghroui et al. 2014). Finally, it can also be noted that all other things being equal, a smaller value of w_j favors the selection of variable y_j in the pricing problem. Conversely, an infinite weight is equivalent to discarding a variable in a partial pricing strategy.

Let $y_V^k := [y_j^k]_{j \in V_p}$ denote an optimal solution to the primal-dual system of value μ_V^k at iteration $k \geq 0$. If $\mu_V^k \geq 0$, the dual optimality conditions are satisfied and we terminate with an optimal solution x^k . Otherwise, $\mu_V^k < 0$ and the current solution might be *improved* by following a direction.

3.4. Projected and Lifted Cycles

An optimal negative cycle, say W_V^k , derived from (21) is augmenting since only residual variables $y_j, j \in V_p$, are considered (see Definition 3). However, the fact that the solution y_V^k for the visible variables is built omitting the variables in $y_H \geq 0$ means that it only provides a *portion* of y^k . Then again, by construction of the linear transformation, these hidden components are uniquely determined within (17) in the system of row set R where the free nature of variables in P is assumed, that is,

$$\sum_{j \in H_p} k_{Rj} y_j + \sum_{j \in V_p} k_{Rj} y_j^k = 0, \quad y_j \geq 0, \quad \forall j \in H_p. \quad (22)$$

The nonsingularity of A_{RP} is made apparent in the transformed system (13) and also finds its way in the equivalent formulation (17). The solution y_H^k to (22) is then determined alike (10) by $\lambda_p := A_{RP}^{-1}(-\sum_{j \in V_p} k_{Rj} y_j^k)$ as

$$y_j^k = \begin{cases} -\lambda_j, & \text{if } \lambda_j < 0 \\ 0, & \text{otherwise} \end{cases} \quad \text{and} \\ y_{j+n}^k = \begin{cases} \lambda_j, & \text{if } \lambda_j > 0 \\ 0, & \text{otherwise} \end{cases} \quad \text{for all } j \in P, \quad (23)$$

and the direction \bar{v}^k follows from the application of Definition 1 to $y^k = [y_H^k, y_V^k, y_Z^k]$. Observe that the complementarity condition $y_j y_{j+n} = 0, \forall j \in \{1, \dots, n\}$, is

taken into account at every stage. First off, by looking for extreme point solutions to the pricing problem, any negative cycle W_V^k cannot contain both y_j and y_{j+n} variables simultaneously. Secondly, y_H^k is established in (23) by dichotomy on the signs of λ_p .

All in all, the cycle W_V^k found in the pricing problem is the support of incomplete information about the direction yet, once W_V^k is identified, the complete cycle is always uniquely determined. In this respect, let W_V^k be called a *projected cycle* on the visible variables in the linear system (21) whereas the full cycle produced with $[y_H^k, y_V^k]$ is named the *lifted cycle* and is denoted by $W_{V \cup H}^k$.

Furthermore, since the reduced cost of the hidden variables is $\bar{d}_j = 0, \forall j \in H_p$, we must have $\bar{d}(W_V^k) = \bar{d}(W_{V \cup H}^k)$. By Lemma 1, the reduced cost of the cycle $W_{V \cup H}^k$ also corresponds to its cost such that fixing the reduced costs of the hidden components to zero transfers the reduced cost information to the visible variables as $\bar{d}_j, j \in V_p$, although the latter contains some dual variables free to be optimized in the pricing problem. Finally, recall that the projected cycle W_V^k is always augmenting. Whether or not the lifted cycle $W_{V \cup H}^k$ is itself guaranteed to be augmenting over the set of residual variables in J^k is directly related to the free nature of the hidden variables in H_p . In other words, whether the step size computed next is certainly positive depends on the content of P (see Proposition 2).

3.5. Step Size and Updates

The step size ρ^k of the lifted cycle $W_{V \cup H}^k$ is computed with respect to the residual capacities of the variables forming it, divided by their respective contribution as

$$\rho^k := \min_{j \in W_{V \cup H}^k} \left\{ \frac{r_j^k}{y_j^k} \right\} \geq 0. \quad (24)$$

A new primal solution $x^{k+1} := x^k + \rho^k \bar{v}^k$ with cost $z^{k+1} := z^k + \rho^k \mu_V^k$ is obtained as

$$\forall j \in \{1, \dots, n\}, \quad x_j^{k+1} := \begin{cases} x_j^k + \rho^k y_j^k, & \text{if } j \in W_{V \cup H}^k \\ x_j^k - \rho^k y_j^k, & \text{if } j+n \in W_{V \cup H}^k \\ x_j^k, & \text{otherwise.} \end{cases} \quad (25)$$

Depending on the choice of the subspace basis A_p , x^{k+1} represented by $[x_F^{k+1}, x_L^{k+1}, x_U^{k+1}]$ could be nonbasic. Section 4.2 explains how and when this can happen with the conceptualization of *interior directions*. We simply mention that any nonbasic solution x^{k+1} can be rendered basic by solving a restricted problem over the set of free variables:

$$z^{k+1} = \min c_F^T x_F + c_L^T x_L + c_U^T x_U \\ \text{s.t. } A_F x_F + A_L x_L + A_U x_U = b \\ l_F \leq x_F \leq u_F, \quad x_L = x_L^{k+1}, \quad x_U = x_U^{k+1}. \quad (26)$$

This problem identifies augmenting cycles comprising free variables only, and increases or decreases the value of these variables until some lower and upper bounds are reached while possibly improving the overall solution cost. In network flows terminology, one obtains a *cycle free solution*, that is, a network solution containing no cycle composed of free arcs only; see Ahuja et al. (1993).

There only remains to update the residual problem $LP(x^{k+1})$ with residual capacities r^{k+1} and column partition $\{F, L, U\}$, and to select a new subspace basis A_P . Another iteration $k \rightarrow k + 1$ then starts in Step 1.

4. Properties

The generic framework ultimately depends on a single parameter, that is, the selection of the set P . Section 4 derives two propositions revolving around this selection. In Section 4.1, we underline particular well-known variants of this framework, whereas Section 4.2 qualifies the kinds of directions found in accordance with the selected set P . In Section 4.3, an illustrative example on a three-dimensional polyhedron shows that a direction with a positive step size can occur on an edge, or be interior. Finally, Section 4.4 draws the line between the proposed framework and actual implementations based on its principles.

4.1. Special Cases

Let us start with a family of variants that perform a positive step size at every iteration. This is completed with four specific variants found in the linear programming literature.

Proposition 2. *Let $x^k, k \geq 0$, be a nonoptimal basic solution to LP (1). Given $P \subseteq F$, the step size of \vec{v}^k is guaranteed to be positive.*

Proof. If $P \subseteq F$, or equivalently $P \cap (L \cup U) = \emptyset$, then all the hidden y -variables in H_P are residual as well as those in V_P . Therefore, the primal-dual pair of the pricing problem (20)–(21), respectively, matches the necessary and sufficient primal-dual optimality conditions of Proposition 1. Indeed, the a posteriori lifted cycle $W_{V \cup H}^k$ obtained from W_V^k trivially only uses variables in J^k and is as such a negative cycle in $LP(x^k)$, meaning that the associated step size is positive. \square

Remark. Consider the case $P \not\subseteq F$, or equivalently $P \cap (L \cup U) \neq \emptyset$, from the perspective of dual optimality conditions. If there exists a $j \in P$ such that $j \notin F$ (i.e., $j \in L \cup U$), then the reduced cost of zero imposed on both, y_j and y_{j+n} is too stringent. The reduced cost of a y -variable with a null residual capacity is irrelevant, which would incidentally have granted more freedom to ψ_R . One can see the pricing oracle construction also misleads the verification of the dual condition (6). Indeed, (18) implies $\bar{d}_j = \bar{d}_{j+n} = 0, \forall j \in P$, or

equivalently $\bar{d}_j = d_j - \psi_R^T k_{R_j} \geq 0, \forall j \in H_P$, that is, the nonfree variables contained in H_P are all part of the verification when they should not. In other words, this overly restrictive observation is also echoed in the primal form where a cycle using such a y -variable is made possible, that is, the additional column in the primal form comes from the additional constraint in the dual form. Observe that this is portrayed in PS when facing a degenerate basis.

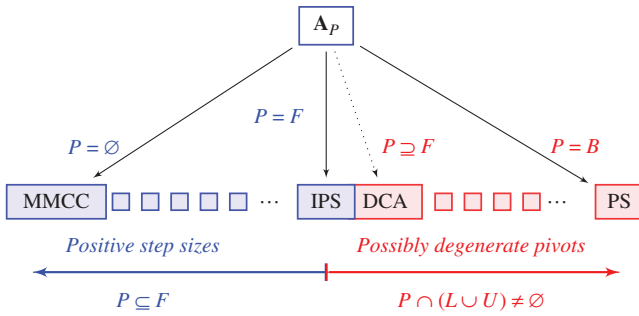
Case $P = \emptyset$. When choosing $P = \emptyset$, it amounts to a subspace basis A_\emptyset of dimension zero, which in turn means that $V(A_\emptyset) = \{\mathbf{0}\}$. Since the vector subspace contains only the zero vector, there are no compatible variables, basic or otherwise. The transformation is trivial, that is, $T_\emptyset = T_\emptyset^{-1} = \begin{bmatrix} I_0 & \emptyset \\ \emptyset & I_m \end{bmatrix} = I_m$. From a dual point of view, the entire m -dimensional dual vector π is optimized to maximize the minimum reduced cost. From a primal point of view, the pricing problem contains precisely all the residual variables and guarantees a positive step size (otherwise the current solution is optimal). When A is the node-arc incidence matrix of a network, this particular case corresponds to the remarkable strongly polynomial minimum mean cycle-canceling algorithm (MMCC) of Goldberg and Tarjan (1989) devised for capacitated minimum cost flow problems. With respect to arbitrary linear programs, it appears natural to think of yet another analogy: *minimum weighted cycle-canceling algorithm*. From a mechanical point of view, the adaptation is straightforward. However, the extent to which time complexity properties are transferable is left for another paper.

Case $P = F$. Choosing $P = F$ corresponds to the strategy developed by Elhallaoui et al. (2011) in IPS. The vector subspace $V(A_F)$ includes all columns of $A_F = \begin{bmatrix} A_{RF} \\ A_{SF} \end{bmatrix}$ but none associated with the basic variables at one of their bounds: $\forall j \in B, a_j \in V(A_F) \Leftrightarrow j \in F$. The transformation is given by $T_F^{-1} = \begin{bmatrix} I_f & 0 \\ -M & I_{m-f} \end{bmatrix}$, where $M = A_{SF}A_{RF}^{-1}$. As a special case of Proposition 2, a positive step size is guaranteed.

Case $P = B$. When choosing $P = B$, we have m linearly independent column vectors with A_B and $V(A_B) = \mathbb{R}^m$. All variables are compatible whereas the sets P and Q , respectively, correspond to all basic and nonbasic variables. The subspace basis induces $T_B = T_B^{-1} = \begin{bmatrix} I_m & \emptyset \\ \emptyset & I_0 \end{bmatrix} = I_m$ yielding once again a trivial transformation. Most importantly, it fixes $\pi^T = c_B^T A_B^{-1}$, that is, all dual variables, and with $w = \mathbf{1}$, the generic framework becomes the primal simplex algorithm with Dantzig's pivot-selection rule. When $B \cap (L \cup U) \neq \emptyset$, that is, when at least one basic degenerate variable is present, the set H_P contains y -variables with null residual capacities and a null step size can occur, that is, a degenerate pivot.

Case $P \supseteq F$. When choosing $P \supseteq F$, we have $f \leq p \leq m$ linearly independent column vectors in A_P . This case

Figure 6. (Color online) Special Cases of VSD: MMCC (1989), IPS (2008), DCA (2005), and PS (1947)



notably includes DCA, the strategy used in Elhallaoui et al. (2005) for solving set partitioning problems by column generation. While the equivalent form with the columns of A_B exists, the subspace basis $\begin{bmatrix} I \\ M \end{bmatrix}$ is obtained by design, more precisely heuristic row clustering as seen in Figure 4. If $P \supset F$, then the set H_P contains $p - f$ y -variables with null residual capacities. This possibly larger than necessary subspace basis gives a lot of freedom in the implementation of DCA, a method steered by practical imperatives.

Figure 6 synthesizes these special cases with respect to Proposition 2, providing also the year they have been designed.

4.2. Interior Directions

Since PS relies on the edge movement induced by a pivot by considering the direction of travel, let us add a layer of definition on the resulting impact of this direction.

Definition 5. Let C be a polyhedron. Given a vertex $x \in C$ and a direction $\vec{v} \neq 0$, let $x + \rho \vec{v} \in C$ for some $\rho > 0$. The vector \vec{v} is called an *edge direction* originating from x if for $0 < \delta < \rho$, the vector $x + \delta \vec{v}$ belongs to an edge of C . Otherwise, a nonedge direction is called an *interior direction* originating from x .

An important property of IPS (where $P = F$) is its movement on an edge (Elhallaoui et al. 2011, Proposition 4). Proposition 2 shows that the family of algorithms with $P \cap (L \cup U) = \emptyset$, or equivalently $P \subseteq F$, ensures a positive step size at every iteration. This also means that the oracle associated with any of these variants is able to verify the necessary and sufficient optimality conditions. While one might rest uneasy about equivalent necessary and sufficient optimality conditions provided by two different oracles, the following proposition sheds light on their content and characterizes improving interior directions originating from a nonoptimal basic solution x^k . In a nut shell, since the case $P = F$ only identifies edge directions, it provides the smallest dimensional cone $\bar{K}_{SV} y_V = 0, y \geq 0$, able to exhaustively identify the set of feasible edge directions. Variants using $P \subset F$ must then contain these

edge directions or combinations of these, that is, interior directions.

Proposition 3. Let $x^k, k \geq 0$, be a nonoptimal basic solution to LP (1). For $P \subset F$, an interior direction \vec{v}^k (if any) is a nonnegative combination of the edge directions.

Proof. For $P = F$, any lifted solution $y^k = [y_V^k, y_H^k, y_Z^k]$ to (21)–(23) is in a one-to-one correspondence with an extreme ray of the cone defined by removing the normalization constraint. Let Ω_F^k be the index set of these extreme rays, indexed by ω . Any solution y to (21)–(23) is nonnull and by the representation theorems of Minkowski and Weyl (see Section 7.2 in Schrijver 1986), it can then be expressed as a nonnegative combination of these extreme rays, that is, $y = \sum_{\omega \in \Omega_F^k} y^\omega \lambda^\omega, \lambda^\omega \geq 0, \forall \omega \in \Omega_F^k$, or component-wise for $j \in \{1, \dots, 2n\}$, $y_j = \sum_{\omega \in \Omega_F^k} y_j^\omega \lambda^\omega, \lambda^\omega \geq 0, \forall \omega \in \Omega_F^k$. By Definition 1, every component $\vec{v}_j, j \in \{1, \dots, n\}$, of the associated direction \vec{v} is given by

$$\vec{v}_j = y_j - y_{j+n} = \sum_{\omega \in \Omega_F^k} (y_j^\omega - y_{j+n}^\omega) \lambda^\omega, \quad \lambda^\omega \geq 0, \forall \omega \in \Omega_F^k,$$

and hence, \vec{v} is a nonnegative combination of the edge directions:

$$\vec{v} = \sum_{\omega \in \Omega_F^k} v^\omega \lambda^\omega, \quad \lambda^\omega \geq 0, \forall \omega \in \Omega_F^k. \quad (27)$$

For $P \subset F$, the pricing problem involves more visible y -variables compared to the case with $P = F$ because $V_P \supset V_F$. At the same time, it contains $f - p$ more constraints since $m - p + 1 > m - f + 1$. Therefore, if \vec{v}^k is not an edge direction, then it is interior and by (27) it can be expressed as a nonnegative combination of the edge directions. \square

Note that a direction leading to a nonbasic solution must be an interior direction. As such, observe that the case $P = \emptyset$ is one of the variants susceptible to lead to nonbasic solutions. However, since $T_\emptyset^{-1} = I_m$ and all dual variables are optimized in the pricing problem, the simplifications applicable to the different steps in Figure 5 imply that maintaining the basic nature of encountered solutions is irrelevant. For all other cases, fetching an index set P of linearly independent columns can be made simply by solving (26).

4.3. Illustrative Examples

Consider the linear program expressed in (28) with x_1, x_2 , and x_3 , and four inequality constraints. Let x_4, \dots, x_7 be the slack variables associated with each constraint and assume the initial basic solution x^0 uses these at values $x_4^0 = 21, x_5^0 = 8, x_6^0 = 15$ and $x_7^0 = 32$ for a total cost of $z^0 = 0$. This basic solution is nondegenerate

Figure 7. Extreme Rays y_V at x^0 in Pricing for $P = F$ and $P = \emptyset$

| Variants | y_1 | y_2 | y_3 | y_4 | y_5 | y_6 | y_7 | y_8 | y_9 | y_{10} | y_{11} | y_{12} | y_{13} | y_{14} |
|-----------------|-------|-------|-------|-------|-------|-------|-------|-------|-------|----------|----------|----------|----------|----------|
| $P = F$ | 1 | 0 | 0 | | | | | | | | | | | |
| | 0 | 1 | 0 | | | | | | | | | | | |
| | 0 | 0 | 1 | | | | | | | | | | | |
| $P = \emptyset$ | 1 | 0 | 0 | 0 | 1 | 0 | 1 | | | | 2 | 0 | 2 | 0 |
| | 0 | 1 | 0 | 1 | 0 | 1 | 1 | | | | 0 | 1 | 0 | 0 |
| | 0 | 0 | 1 | 0 | 1 | 1 | 0 | | | | 2 | 0 | 0 | 2 |
| | 0 | 2 | 1 | 0 | 0 | 3 | 0 | | | | 0 | 1 | 0 | 0 |
| | 1 | 2 | 0 | 0 | 0 | 0 | 3 | | | | 0 | 1 | 0 | 0 |
| | 2 | 3 | 1 | 0 | 0 | 0 | 3 | | | | 3 | 0 | 0 | 0 |

and hence there are three edge directions according to the selected entering variable $x_1, x_2,$ or x_3 .

$$\begin{aligned}
 \max \quad & 130x_1 + 80x_2 + 60x_3 \\
 \text{s.t.} \quad & 2x_1 - x_2 + 2x_3 \leq 21 \\
 & -x_1 + x_2 - x_3 \leq 8 \\
 & 2x_1 - x_2 - x_3 \leq 15 \\
 & -x_1 - x_2 + 2x_3 \leq 32 \\
 & x_1, x_2, x_3 \geq 0.
 \end{aligned} \tag{28}$$


Let us start with the possible values of the projected vector $y_V \in \mathbb{R}_+^{|V|}$ arising from $y \in \mathbb{R}_+^{14}$ while solving the oracle (21). For $P = F$, eight variables are hidden from the pricing problem, that is, the forward and backward y -variables associated with the four slack variables in F , only three are visible in (21), that is, y_1, y_2 and y_3 , whereas y_8, y_9 and y_{10} are fixed to zero. The three-dimensional extreme ray (y_1, y_2, y_3) can take values $(1, 0, 0), (0, 1, 0)$, and $(0, 0, 1)$. For $P = \emptyset$ and incidentally $H_P = \emptyset$, y_8, y_9 and y_{10} are again fixed to zero such that all the residual variables are visible in the oracle. The vector y_V is therefore the 11-dimensional vector $(y_1, \dots, y_7, y_{11}, \dots, y_{14})$ for which the six extreme rays, conveniently expressed with integers, are listed in Figure 7.

Figure 8 details the content of $\vec{v}^0 \in \mathbb{R}^7$ found under the suggested set P and weight vector w . One would of course use the stopping criterion $\mu_V^k \leq 0$ since (28) is being maximized. The minimum cost of the associated cycle W_V^0 is given by μ_V^0 , the step size ρ^0 recovered

Figure 8. Directions \vec{v}^0 Found at x^0 in Pricing for $P = F$ and Two Variants of $P = \emptyset$

| Variant | w_j | c_j coefficients variables | 130 | 80 | 60 | | | | | | μ_V^0 | ρ^0 |
|-----------------|-------------|--------------------------------------|---------------------|----------------------|---------------------|---------------------|----------------------|-------|-------|----------------------|-----------|-------------------------|
| | | | x_1 | x_2 | x_3 | x_4 | x_5 | x_6 | x_7 | z^1 | | |
| $P = F$ | 1 | \vec{v}^0 (edge) x^1 | 1 | 0 | 0 | -2 | 1 | -2 | 1 | | 130 | 7.5 |
| $P = \emptyset$ | 1 | \vec{v}^0 (interior) x^1 | $\frac{1}{6}$ 14 | $\frac{1}{4}$ 21 | $\frac{1}{12}$ 7 | $-\frac{1}{4}$ 0 | 0 | 8 | 15 | 53 | 975 | $\frac{140}{3}$ 84 |
| $P = \emptyset$ | $\ a_j\ ^2$ | \vec{v}^0 (interior/face) x^1 | $\frac{1}{22}$ 8 | $\frac{1}{11}$ 16 | 0 | 0 | $-\frac{1}{22}$ 0 | 0 | 15 | $\frac{3}{22}$ 56 | 2,320 | $\frac{145}{11}$ 176 |

Figure 9. (Color online) Three-Dimensional Interior Direction $(\vec{v}_1^0, \vec{v}_2^0, \vec{v}_3^0) = (\frac{1}{6}, \frac{1}{4}, \frac{1}{12})$ with $P = \emptyset$ and $w = \mathbf{1}$



with (24), and the new solution x^1 obtained with (25). As expected, the direction $\vec{v}^0 = (1, 0, 0, -2, 1, -2, 1)$ found with IPS ($P = F$) follows an edge and yields an extreme point solution at x^1 . Notice that since x^0 is non-degenerate, we have $F = B$ and we would have found the same direction upon selecting x_1 as the entering variable in PS.

When $P = \emptyset$ and $w_j = 1, \forall j \in \{1, \dots, 14\}$, $\vec{v}^0 = (\frac{1}{6}, \frac{1}{4}, \frac{1}{12}, -\frac{1}{4}, 0, 0, \frac{1}{4})$ induced by the last extreme ray of Figure 7 so happens to be interior (Figure 9) and yields a feasible solution x^1 , which is nonbasic (six variables take positive values). By Proposition 3, this \vec{v}^0 is indeed the combination of edge directions, that is,

$$\begin{aligned}
 \vec{v}^0 = & \frac{1}{6} (1, 0, 0, -2, 1, -2, 1) + \frac{1}{4} (0, 1, 0, 1, -1, 1, 1) \\
 & + \frac{1}{12} (0, 0, 1, -2, 1, 1, -2).
 \end{aligned}$$

The last example in Figure 8 still uses the set $P = \emptyset$ but uses a weight vector whose every element w_j is determined by computing the squared norm $\|a_j\|^2$ of each column, that is, $w_1 = 2^2 + (-1)^2 + 2^2 + (-1)^2 = 10$. The pricing problem finds a different extreme ray that also induces an interior direction, indeed $\vec{v}^0 = (\frac{1}{22}, \frac{1}{11}, 0, 0, -\frac{1}{22}, 0, \frac{3}{22})$ that happens to be within the x_1x_2 -face.

By no means do we imply that the set $P = \emptyset$ provides all around better directions than with $P = F$. In fact, it suffices to modify the coefficients of x_1 and x_2 in the third constraint to 1 and 3 to get the opposite effect when using $w = \mathbf{1}$. In other words, it remains difficult to distinguish what the quality of a direction really is. For instance, Klee and Minty (1972) propose

a parametric LP for which the Dantzig's pivot rule behaves exponentially. This is a pathological example not reflecting typical practical behavior but yielding a theoretical insight. Similar examples have been suggested for many known pivoting rules (Paparrizos et al. 2009), and we do so within our framework in the following.

The Klee–Minty polytope is a hypercube of parametric dimension whose corners have been distorted. It can be written as

$$\begin{aligned} \max_{x \geq 0} \quad & 2^{n-1}x_1 + 2^{n-2}x_2 + \cdots + 2x_{n-1} + x_n \\ \text{s.t.} \quad & x_1 \leq 5 \\ & 4x_1 + x_2 \leq 25 \\ & 8x_1 + 4x_2 + x_3 \leq 125 \\ & \vdots \\ & 2^n x_1 + 2^{n-1}x_2 + \cdots + 4x_{n-1} + x_n \leq 5^n. \end{aligned} \quad (29)$$

The previous LP contains n variables, n constraints, and 2^n extreme points. Starting at $x = 0$, the initial basis comprises all the slack variables s_j , for $j = 1, \dots, n$, valued at $s_j = 5^j$. Using the Dantzig's pivot rule, PS first selects the variable with the largest reduced cost (indeed, the largest c_j in this case), that is, x_1 as the entering variable. It then goes through each of the extreme points before reaching the optimal solution at $(0, 0, \dots, 5^n)$. None of the $2^n - 1$ pivots is degenerate, hence IPS would follow the same trajectory as PS, this showing that IPS can also have an exponential time complexity in the worst case. MWCC behaves in a totally different way. Assuming unit weights in the pricing problem (21), the latter writes as

$$\begin{aligned} \mu = \max_{y, s \geq 0} \quad & 2^{n-1}y_1 + 2^{n-2}y_2 + \cdots + 2y_{n-1} + y_n \\ \text{s.t.} \quad & y_1 - y_{1+3n} = 0 \\ & 4y_1 + y_2 - y_{2+3n} = 0 \\ & 8y_1 + 4y_2 + y_3 - y_{3+3n} = 0 \\ & \vdots \\ & 2^n y_1 + 2^{n-1}y_2 + \cdots + 4y_{n-1} + y_n - y_{4n} = 0 \\ & y_1 + y_2 + \cdots + y_{n-1} + y_n + \sum_{j=1}^n y_{j+3n} = 1, \end{aligned} \quad (30)$$

where $y_{1+3n}, y_{2+3n}, \dots, y_{4n}$ are the backward residual variables for the slacks at their upper bounds. Taking into account the last equality to zero together with the normalizing constraint, we derive an upper bound on the value of μ :

$$\begin{aligned} \mu &= \max(2^{n-1}y_1 + 2^{n-2}y_2 + \cdots + 2y_{n-1}) + y_n \\ &= \max\left(\frac{y_{4n} - y_n}{2}\right) + y_n = \max\frac{y_{4n} + y_n}{2} \leq \frac{1}{2}. \end{aligned}$$

Let $\mu(y_j) = 2^{n-j}/(2(2^{n-j+1} - 1))$, $j = 1, \dots, n$, be the value of the objective function in (30) for the cycle using variable y_j , each one inducing an edge direction. The reader can verify that $\mu(y_n) = \frac{1}{2}$, whereas $\mu(y_j) < \frac{1}{2}$ for all $j = 1, \dots, n - 1$. Therefore, MWCC selects the cycle composed of $y_n = 0.5$ and $y_{4n} = 0.5$. The pivot operation raises x_n to 5^n , decreases s_n by the same amount to zero, and therefore MWCC reaches the optimal solution in a single iteration on the Klee–Minty polytope.

4.4. Computational Behavior

Having a generic framework that conceptually hosts, for example, PS, IPS, DCA, and MMCC does not imply that we believe in a generic competitive implementation of the framework. Note that all these specializations are efficiently implemented in the literature, but each implementation massively exploits their respective context. However, our unified presentation opens up further opportunities for more tailoring to particular applications. This may manifest in special kinds of transformation matrices, special update strategies of the row/column partitions, and/or in special strategies of computing dual solutions. We think of this versatility as one of the framework's biggest strengths that may be most useful in contexts in which solving the pricing problem is expensive anyway, like in column generation. An oracle capable of producing strictly improving columns may be extremely useful.

PS at one end of the spectrum is highly competitive, even in network problems where 70%–90% of the pivots are degenerate (Ahuja et al. 1993, Section 18.3). At the other end, MMCC applied on the same type of problems is desperately slow where sometimes solving the oracle takes more time than the original problem (Kovács 2015); we observe the same behavior for the linear programming counterpart MWCC. Nevertheless, when cancel-and-tighten, a clever partial pricing strategy for which the order of the improving cycles becomes irrelevant, is considered alongside MMCC, we see significant time reductions (Gauthier et al. 2015, 2017). In this light, MWCC, just like its network counterpart MMCC, may be an intermediate step to reach a more efficient cancel-and-tighten version adapted to linear programs. This is certainly worth investigating.

Moreover, while IPS requires the inverse of a submatrix within the current basis, the additional computational burden needs to be intimately tied to the implementation for it to make any sense. On that note, it seems that the *positive edge* rule is already implemented in practice (Towhidi et al. 2014).

Finally, DCA is a sound and proven method based on an intuitive clustering of the rows appearing in master problems based on set partitioning formulations. Multiple strategies have been developed to cope with routing and scheduling applications in an efficient manner (Elhallaoui et al. 2008, 2010). While it does not seem

like it can efficiently be applied to arbitrary linear programming matrices, DCA has also been extended to ISUD, which considers integrality in the solution process. The idea is to transfer the integrality requirements to the oracle. The latter aims to determine integer leading directions via ad hoc cutting planes and dynamic weights for the normalization constraint until it fails for lack of tractability. ISUD is as such often able to terminate with a near-optimal solution thus eliminating the typical branch-and-bound tree.

5. Conclusion

This paper unites under one generic framework several known primal algorithms with a broad spectrum of possibilities. Aside from pinpointing reasons simplex-type algorithms suffer from degeneracy, the elimination of the latter is made possible through a linear transformation. The purpose of this paper is further driven by primal algorithms such as column generation where mechanisms coping with degeneracy are beneficial. The two extreme cases of our framework correspond to the primal simplex and the minimum weighted cycle-canceling algorithms. Two properties are established for different family members: positive step sizes at every iteration and pricing problems that provide edge directions only. The improved primal simplex algorithm is remarkably the only variant that qualifies for both features. While interior directions are certainly usual in the realm of nonlinear algorithms, it is not so often that one thinks about such possibilities for simplex-like algorithms. On another note, the minimum weighted cycle-canceling algorithm requires neither a matrix transformation nor the maintenance of basic solutions. It even does not require the knowledge of a dual vector at any iteration, although the convergence proof in strongly polynomial time for the network version is totally driven by the dual point of view. Because the cost and the reduced cost of a cycle are equal, any heuristic looking for negative cycles could interchangeably use the original costs of the variables or some reduced costs derived from any approximation of the dual values. This idea is already proving elegant in our recent experiments with heuristic subproblem pricers within branch and price for vehicle routing problems where columns are found using the original costs rather than the traditional reduced costs.

The vector space decomposition framework revolves around a unique parameter and is derived directly from necessary and sufficient optimality conditions established on the residual problem. The parameter may vary at every iteration and dictates how the linear program decomposition is made. Some variables are hidden, producing an oracle looking for cycles in a projected space where only visible variables remain.

Cycles are found and lifted back to the residual problem before finding the step size along the associated direction.

Finally, column generation as a primal algorithm to solve large-scale linear programs is a main beneficiary of our proposal. Indeed, the *dual guided* pricing can reduce the row size of the master problem where degeneracy difficulties occur. Moreover, accelerating strategies such as stabilization techniques can be incorporated to all the variants. Indeed, dual variables can be optimized within intervals in either the master or the subproblem providing flexible arrangements. As we play with dual variables, the automated identification of dual optimal inequalities (Valério de Carvalho 2005, Ben Amor et al. 2006, Gschwind and Irnich 2016) is also appealing.

References

- Ahuja RK, Magnanti TL, Orlin JB (1993) *Network Flows: Theory, Algorithms, and Applications* (Prentice Hall, Upper Saddle River, NJ).
- Bazaraa MS, Jarvis JJ, Sherali HD (1990) *Linear Programming and Network Flows* (John Wiley & Sons, New York).
- Ben Amor HMT, Desrosiers J, Frangioni A (2009) On the choice of explicit stabilizing terms in column generation. *Discrete Appl. Math.* 157(6):1167–1184.
- Ben Amor HMT, Desrosiers J, Valério de Carvalho JM (2006) Dual-optimal inequalities for stabilized column generation. *Oper. Res.* 54(3):454–463.
- Bixby RE (2002) Solving real-world linear programs: A decade and more of progress. *Oper. Res.* 50(1):3–15.
- Dantzig GB, Thapa MN (1997) *Linear Programming 1: Introduction*. Mikosch TV, Resnick SI, Robinson SM, eds. Springer Series in Operations Research and Financial Engineering (Springer, New York).
- Dantzig GB, Thapa MN (2003) *Linear Programming 2: Theory and Extensions*. Mikosch TV, Resnick SI, Robinson SM, eds. Springer Series in Operations Research and Financial Engineering (Springer, New York).
- Dantzig GB, Wolfe P (1960) Decomposition principle for linear programs. *Oper. Res.* 8(1):101–111.
- du Merle O, Villeneuve D, Desrosiers J, Hansen P (1999) Stabilized column generation. *Discrete Math.* 194(1–3):229–237.
- Elhallaoui I, Desaulniers G, Metrane A, Soumis F (2008) Bi-dynamic constraint aggregation and subproblem reduction. *Comput. Oper. Res.* 35(5):1713–1724.
- Elhallaoui I, Metrane A, Desaulniers G, Soumis F (2011) An improved primal simplex algorithm for degenerate linear programs. *INFORMS J. Comput.* 23(4):569–577.
- Elhallaoui I, Metrane A, Soumis F, Desaulniers G (2010) Multi-phase dynamic constraint aggregation for set partitioning type problems. *Math. Programming* 123(2):345–370.
- Elhallaoui I, Villeneuve D, Soumis F, Desaulniers G (2005) Dynamic aggregation of set partitioning constraints in column generation. *Oper. Res.* 53(4):632–645.
- Forrest JJ, Goldfarb D (1992) Steepest-edge simplex algorithms for linear programming. *Math. Programming* 57(1):341–374.
- Gauthier JB, Desrosiers J, Lübbecke ME (2014) Decomposition theorems for linear programs. *Oper. Res. Lett.* 42(8):553–557.
- Gauthier JB, Desrosiers J, Lübbecke ME (2015) About the minimum mean cycle-canceling algorithm. *Discrete Appl. Math.* 196(December):115–134.
- Gauthier JB, Desrosiers J, Lübbecke ME (2016) Tools for primal degenerate linear programs: IPS, DCA, and PE. *EURO J. Transportation Logist.* 5(2):161–204.
- Gauthier JB, Desrosiers J, Lübbecke ME (2017) A strongly polynomial contraction-expansion algorithm for network flow problems. *Comput. Oper. Res.* 84(August):16–32.

- Goldberg AV, Tarjan RE (1989) Finding minimum-cost circulations by canceling negative cycles. *J. ACM* 36(4):873–886.
- Gschwind T, Irnich S (2016) Dual inequalities for stabilized column generation revisited. *INFORMS J. Comput.* 28(1):175–194.
- Harris PMJ (1973) Pivot selection methods of the Devex LP code. *Math. Programming* 5(1):1–28.
- Klee V, Minty GJ (1972) How good is the simplex algorithm? Shisha O, ed. *Inequalities*, Vol. 3 (Academic Press, New York), 159–175.
- Kovács P (2015) Minimum-cost flow algorithms: An experimental evaluation. *Optim. Methods Software* 30(1):94–127.
- Lübbecke ME, Desrosiers J (2005) Selected topics in column generation. *Oper. Res.* 53(6):1007–1023.
- Paparrizos K, Samaras N, Zissopoulos D (2009) Linear programming: Klee–Minty examples. Floudas CA, Pardalos PM, eds. *Encyclopedia of Optimization*, 2nd ed. (Springer, Boston), 1891–1897.
- Rosat S, Elhallaoui I, Soumis F, Chakour D (2017a) Influence of the normalization constraint on the integral simplex using decomposition. *Discrete Appl. Math.* 217(January):53–70.
- Rosat S, Quesnel F, El Hallaoui I, Soumis F (2017b) Dynamic penalization of fractional directions in the integral simplex using decomposition: Application to aircrew scheduling. *Eur. J. Oper. Res.* 263(3):1007–1018.
- Schrijver A (1986) *Theory of Linear and Integer Programming* (John Wiley & Sons, Chichester, West Sussex, UK).
- Towhidi M, Desrosiers J, Soumis F (2014) The positive edge criterion within COIN-OR’s CLP. *Oper. Res.* 49(September):41–46.
- Valério de Carvalho JM (2005) Using extra dual cuts to accelerate convergence in column generation. *INFORMS J. Comput.* 17(2):175–182.
- Zaghroui A, Soumis F, El Hallaoui I (2014) Integral simplex using decomposition for the set partitioning problem. *Oper. Res.* 62(2):435–449.

Jean Bertrand Gauthier is currently at the Johannes Gutenberg University in Mainz, Germany, in a postdoctoral position with Professor Stefan Irnich. His work, “Primal Algorithms for Degenerate Linear and Network Flow Problems,” was awarded the best dissertation prize at HEC Montréal (2016).

Jacques Desrosiers is a full professor in the Department of Management Sciences at HEC Montréal. He is also a member of the GERAD Operations Research Center. His main research interests include column generation algorithms and large-scale optimization for vehicle routing and crew scheduling in air, rail, and urban transportation.

Marco E. Lübbecke is a professor and chair of operations research at the School of Business and Economics and the Department of Mathematics at RWTH Aachen University, Germany. He has been modeling and solving large and complex discrete optimization problems from science and industry for more than 20 years. His main research contributions are in computational integer programming, in particular in the field of decomposition algorithms. Marco is an INFORMS board member.