

Decomposition and Dynamic Cut Generation in Integer Linear Programming

T.K. Ralphs* M.V. Galati†

Revised August 14, 2004

Abstract

Decomposition algorithms such as Lagrangian relaxation and Dantzig-Wolfe decomposition are well-known methods that can be used to generate bounds for mixed-integer linear programming problems. Traditionally, these methods have been viewed as distinct from polyhedral methods, in which bounds are obtained by dynamically generating valid inequalities to strengthen a linear programming relaxation. Recently, a number of authors have proposed methods for integrating dynamic cut generation with various decomposition methods to yield further improvement in the computed bounds. In this paper, we describe a framework within which most of these methods can be viewed from a common theoretical perspective. We then show how the framework can be extended to obtain a new technique we call *decompose and cut*. As a by-product, we describe how these methods can take advantage of the fact that solutions with known structure, such as those to a given relaxation, can frequently be separated much more easily than arbitrary solutions.

1 Introduction

In this paper, we discuss methods for computing bounds on the value of an optimal solution to a mixed-integer linear program (MILP). Computing such bounds is an essential element of the branch and bound algorithm, which is the most effective and most commonly used method for solving mathematical programs. Bounds are generally computed by solving a *bounding subproblem*, which is either a relaxation or a dual of the original problem. The most commonly used bounding subproblem is the linear programming (LP) relaxation. The LP relaxation is often too weak to be effective, but it can be strengthened by the addition of dynamically generated valid inequalities. Alternatively, traditional decomposition techniques, such as Dantzig-Wolfe decomposition [19] or Lagrangian relaxation [22, 14], can also be used to obtain improved bounds.

Methods based on cut generation have traditionally been considered distinct from decomposition methods, but several authors have suggested combining the two approaches

*Department of Industrial and Systems Engineering, Lehigh University, Bethlehem, PA 18015, tkralphs@lehigh.edu, <http://www.lehigh.edu/~tkr2>

†Analytical Solutions - Operations R & D, SAS Institute, Cary, NC 27513 matthew.galati@sas.com, <http://sagan.ie.lehigh.edu/mgalati>

to yield further improvements (these contributions are reviewed in Section 3). In this paper, we present a framework that shows how these hybrid methods can be viewed from a common theoretical perspective as generalizations of the cutting plane method. From this viewpoint, we show how various methods from the literature are related, as well as present a new method called *decompose and cut* that follows the same paradigm. One element common to these methods is the ability to take advantage of the fact that the separation problem is often much easier for solutions with combinatorial structure than for arbitrary real vectors. This can simplify the implementation of these methods and allow more rapid development than with traditional cutting plane implementations. We discuss the importance of this and provide several examples of its usefulness.

The goal of this paper is not to provide a computational comparison of methods. Such comparisons are necessarily problem-dependent and based largely on empirical evidence. Although some general statements can be made, the lines between the various methods are blurry at best and many subtle variations are possible. Our goal here is to provide the reader with insight, which must then be combined with experience to guide the choice of method. By illustrating the relationship between various methods, we also provide a methodological framework within which it is easy to switch between variants by replacing certain algorithmic components. We are currently developing a software framework that will provide just such a capability.

To simplify the exposition, we consider only pure integer linear programming problems (ILPs) with bounded feasible regions. The framework can easily be extended to more general settings. For the remainder of the paper, we consider an ILP whose feasible set is the integer vectors contained in the polyhedron $\mathcal{Q} = \{x \in \mathbb{R}^n : Ax \geq b\}$, where $A \in \mathbb{Q}^{m \times n}$ is the constraint matrix and $b \in \mathbb{Q}^m$ is the right-hand-side vector. Let $\mathcal{F} = \mathcal{Q} \cap \mathbb{Z}^n$ be the feasible set and let \mathcal{P} be the convex hull of \mathcal{F} .

We consider two problems associated with the polyhedron \mathcal{P} . The *optimization problem* for \mathcal{P} is that of determining

$$z_{IP} = \min_{x \in \mathcal{F}} \{c^\top x\} = \min_{x \in \mathcal{P}} \{c^\top x\} = \min_{x \in \mathbb{Z}^n} \{c^\top x : Ax \geq b\} \quad (1)$$

for a given cost vector $c \in \mathbb{Q}^n$, where z_{IP} is infinite if \mathcal{F} is empty. A related problem is the *separation problem* for \mathcal{P} . Given $x \in \mathbb{R}^n$, the problem of separating x from \mathcal{P} is that of deciding whether $x \in \mathcal{P}$ and if not, determining $a \in \mathbb{R}^n$ and $\beta \in \mathbb{R}$ such that $a^\top y \geq \beta \forall y \in \mathcal{P}$ but $a^\top x < \beta$. A pair $(a, \beta) \in \mathbb{R}^{n+1}$ such that $a^\top y \geq \beta \forall y \in \mathcal{P}$ is a *valid inequality* for \mathcal{P} and is said to be *violated* by $x \in \mathbb{R}^n$ if $a^\top x < \beta$. In [30], it was shown that the separation problem for \mathcal{P} is polynomially equivalent to the optimization problem for \mathcal{P} .

To apply the principle of decomposition, we consider the relaxation of (1) defined by

$$\min_{x \in \mathcal{F}'} \{c^\top x\} = \min_{x \in \mathcal{P}'} \{c^\top x\} = \min_{x \in \mathbb{Z}^n} \{c^\top x : A'x \geq b'\}, \quad (2)$$

where $\mathcal{F} \subset \mathcal{F}' = \{x \in \mathbb{Z}^n : A'x \geq b'\}$ for some $A' \in \mathbb{Q}^{m' \times n}$, $b' \in \mathbb{Q}^{m'}$ and \mathcal{P}' is the convex hull of \mathcal{F}' . As usual, we assume that there exists an *effective* algorithm for optimizing over \mathcal{P}' . We are deliberately using the term *effective* here to denote an algorithm that has an acceptable average-case running time, since this is more relevant than worst-case behavior in a computational framework. Along with \mathcal{P}' is associated a set of *side constraints*. Let

$[A'', b''] \in \mathbb{Q}^{m'' \times n}$ be a set of additional inequalities needed to describe \mathcal{F} , i.e., $[A'', b'']$ is such that $\mathcal{F} = \{x \in \mathbb{Z}^n : A'x \geq b', A''x \geq b''\}$. We denote by \mathcal{Q}' the polyhedron described by the inequalities $[A', b']$ and by \mathcal{Q}'' the polyhedron described by the inequalities $[A'', b'']$. Hence, the initial LP relaxation is the linear program defined by $\mathcal{Q} = \mathcal{Q}' \cap \mathcal{Q}''$, and the *LP bound* is given by

$$z_{LP} = \min_{x \in \mathcal{Q}} \{c^\top x\} = \min_{x \in \mathbb{R}^n} \{c^\top x : A'x \geq b', A''x \geq b''\}. \quad (3)$$

Note that this is slightly more general than the traditional framework, in which $[A', b']$ and $[A'', b'']$ are a partition of the rows of $[A, b]$.

2 Traditional Decomposition Methods

The goal of the decomposition approach is to improve on the LP bound by taking advantage of the ability to optimize over and/or separate from \mathcal{P}' . We briefly review the classical bounding methods that take this approach in order to establish terminology and notation.

Lagrangian Relaxation. For a given vector of *dual multipliers* $u \in \mathbb{R}_+^{m''}$, the *Lagrangian relaxation* of (1) is given by

$$z_{LR}(u) = \min_{s \in \mathcal{F}'} \{(c^\top - u^\top A'')s + u^\top b''\} \quad (4)$$

It is easily shown that $z_{LR}(u)$ is a lower bound on z_{IP} for any $u \geq 0$. The problem

$$z_{LD} = \max_{u \in \mathbb{R}_+^{m''}} \{z_{LR}(u)\} \quad (5)$$

of maximizing this bound over all choices of dual multipliers is a dual to (1) called the *Lagrangian dual* (LD) and also provides a lower bound z_{LD} , which we call the *LD bound*. A vector of multipliers that yield the largest bound are called *optimal (dual) multipliers*. For the remainder of the paper, let \hat{u} be such a vector.

The Lagrangian dual can be solved by any of a number of subgradient-based optimization procedures or by rewriting it as the equivalent linear program

$$z_{LD} = \max_{\alpha \in \mathbb{R}, u \in \mathbb{R}_+^{m''}} \{\alpha + u^\top b'' : \alpha \leq (c^\top - u^\top A'')s \forall s \in \mathcal{F}'\}. \quad (6)$$

and solving it using a cutting plane algorithm. In either case, the computational effort goes into evaluating $z_{LR}(u)$ for a given sequence of dual multipliers u . This is an optimization problem over \mathcal{P}' and can be solved effectively by assumption. This general approach is described in detail in [32].

Dantzig-Wolfe Decomposition. The approach of Dantzig-Wolfe decomposition is to reformulate (1) by implicitly requiring the solution to be a member of \mathcal{F}' , while explicitly enforcing the inequalities $[A'', b'']$. Relaxing the integrality constraints of this reformulation, we obtain the linear program

$$z_{DW} = \min_{\lambda \in \mathbb{R}_+^{\mathcal{F}'}} \{c^\top (\sum_{s \in \mathcal{F}'} s \lambda_s) : A'' (\sum_{s \in \mathcal{F}'} s \lambda_s) \geq b'', \sum_{s \in \mathcal{F}'} \lambda_s = 1\}, \quad (7)$$

which we call the *Dantzig-Wolfe LP* (DWLP). Although the number of columns in this linear program is $|\mathcal{F}'|$, it can be solved by dynamic column generation, where the column-generation subproblem is again an optimization problem over \mathcal{P}' .

It is easy to verify that the DWLP is the dual of (6), which immediately shows that $z_{DW} = z_{LD}$ (see [50] for a treatment of this fact). Hence, z_{DW} is a valid lower bound on z_{IP} that we call the *DW bound*. Note that if we let $\hat{\alpha} = z_{LR}(\hat{u}) - \hat{u}^\top b''$, then $(\hat{u}, \hat{\alpha})$ is an optimal solution for the LP (6) and hence also an optimal dual solution to the DWLP. The optimal primal solution to (7), $\hat{\lambda}$, is referred to as an *optimal Dantzig-Wolfe (DW) decomposition*. If we combine the members of \mathcal{F}' using $\hat{\lambda}$, to obtain

$$\hat{x}_{DW} = \sum_{s \in \mathcal{F}'} s \hat{\lambda}_s, \quad (8)$$

then we see that $z_{DW} = c^\top \hat{x}_{DW}$. Since \hat{x}_{DW} must lie within $\mathcal{P}' \subseteq \mathcal{Q}'$ and also within \mathcal{Q}'' , this shows that $z_{DW} \geq z_{LP}$.

Cutting Plane Method. In the cutting plane method, inequalities describing \mathcal{P}' (i.e., the facet-defining inequalities) are generated dynamically by separating the solutions to a series of LP relaxations from \mathcal{P}' . In this way, the initial LP relaxation is iteratively augmented to obtain the *CP bound*,

$$z_{CP} = \min_{x \in \mathcal{P}'} \{c^\top x : A''x \geq b''\}. \quad (9)$$

We refer to this augmented linear program as the *cutting plane LP* (CPLP) and any optimal solution to CPLP as an *optimal fractional solution*. Note that \hat{x}_{DW} , as defined in (8), is an optimal solution to this augmented linear program. Hence, the CP bound is equal to both the DW bound and the LD bound.

A Common Framework. The following well-known result of Geoffrion [26] relates the three methods just described.

Theorem 1 $z_{IP} \geq z_{LD} = z_{DW} = z_{CP} \geq z_{LP}$.

A graphical depiction of this bound is shown in Figure 1. Theorem 1 shows that (5), (7), and (9) represent three different approaches to computing the same quantity. In each case, we are given a polyhedron \mathcal{P} over which we would like to optimize, along with two additional polyhedra, denoted here by \mathcal{Q}'' and \mathcal{P}' , each of which contain \mathcal{P} . The polyhedron \mathcal{Q}'' has a small description that can be represented explicitly, while the polyhedron \mathcal{P}' has a much larger description and is represented implicitly, i.e., portions of the description are generated dynamically using our ability to effectively optimize/separate. To describe their roles in this framework, we call \mathcal{P} the *original polyhedron*, \mathcal{Q}'' the *explicit polyhedron*, and \mathcal{P}' the *implicit polyhedron*.

By the Weyl-Minkowski Theorem, every bounded rational polyhedron has two descriptions – one as the intersection of half-spaces (the *outer representation*) and one as the convex hull of its extreme points (the *inner representation*) [49]. Conceptually, the difference between the representations utilized by the three methods is that Dantzig-Wolfe decomposition

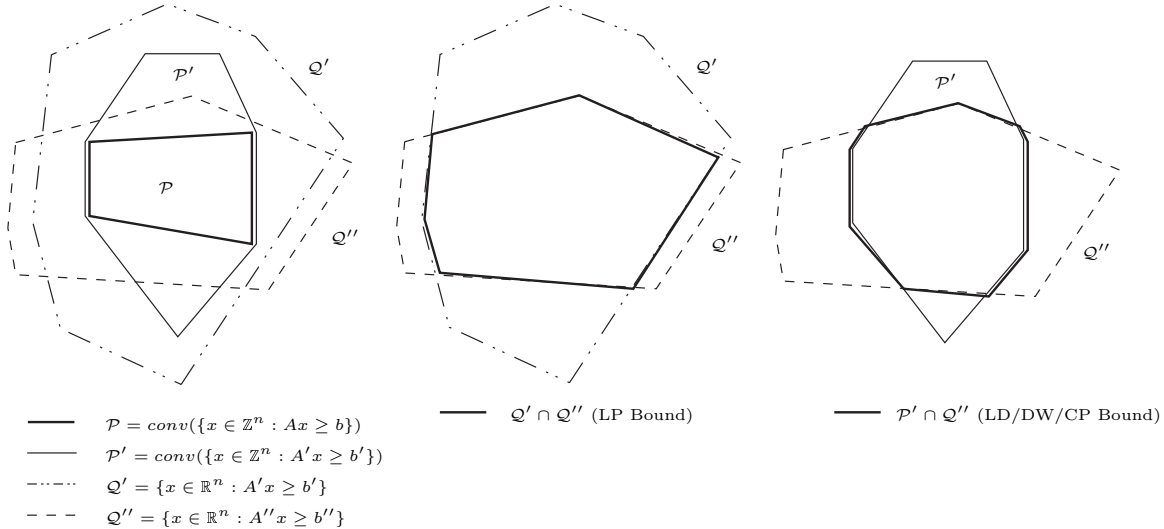


Figure 1: Illustration of the LP and LD/DW/CP bounds

and Lagrangian relaxation utilize an inner representation of \mathcal{P}' , generated dynamically by solving the corresponding optimization problem, whereas the cutting plane method relies on an outer representation of \mathcal{P}' , generated dynamically by solving the separation problem for \mathcal{P}' . For this reason, we will call Dantzig-Wolfe decomposition and Lagrangian relaxation *inner methods*, and the cutting plane method an *outer method*. In theory, we have the same choice of representation for the explicit polyhedron. It is intriguing to ponder the implications of this choice. Note that this framework encompasses methods not normally thought of as decomposition methods. In particular, by defining the implicit polyhedron to be a polyhedron containing \mathcal{P} for which there exists an effective algorithm for solving the separation problem, we can view cutting plane methods in the same frame of reference as we view other decomposition methods. This viewpoint sheds new light on their relationship to traditional decomposition methods.

3 Integrated Decomposition Methods

One of the advantages of outer methods over inner methods is the option of adding heuristically generated valid inequalities to the cutting plane LP to further improve the bound. This can be thought of as a dynamic tightening of either the explicit or the implicit polyhedron. Viewing it as a dynamic tightening of the explicit polyhedron yields a generalization of the cutting plane method obtained by replacing the cutting plane LP with either a Dantzig-Wolfe LP or a Lagrangian dual as the bounding subproblem in a straightforward fashion. We call this class of bounding procedures *integrated decomposition methods* because they integrate inner and outer methods. The steps of this generalized method are shown in Figure 2.

The important step in Figure 2 is Step 3, generating a set of *improving inequalities*, i.e., inequalities valid for \mathcal{P} that when added to the description of the explicit polyhedron result

An Integrated Decomposition Method

Input: An instance of ILP.

Output: A lower bound z on the optimal solution value for the instance.

1. Construct the initial bounding subproblem P^0 and set $i \leftarrow 0$.

$$z_{CP} = \min_{x \in \mathcal{P}'} \{c^\top x : A''x \geq b''\}$$

$$z_{LD} = \max_{u \in \mathbb{R}^{m''}} \min_{s \in \mathcal{F}'} \{(c^\top - u^\top A'')s + u^\top b''\}$$

$$z_{DW} = \min_{\lambda \in \mathbb{R}_+^{\mathcal{F}'}} \{c^\top (\sum_{s \in \mathcal{F}'} s \lambda_s) : A''(\sum_{s \in \mathcal{F}'} s \lambda_s) \geq b'', \sum_{s \in \mathcal{F}'} \lambda_s = 1\}$$

2. Solve P^i to obtain a valid lower bound z^i .
3. Attempt to generate a set of improving inequalities $[D^i, d^i]$ valid for \mathcal{P} .
4. If valid inequalities were found in Step 3, form the bounding subproblem P^{i+1} by setting $[A'', b''] \leftarrow \begin{bmatrix} A'' & b'' \\ D^i & d^i \end{bmatrix}$. Then, set $i \leftarrow i + 1$ and go to Step 2.
5. If no valid inequalities were found in Step 3, then output z^i .

Figure 2: Basic outline of an integrated decomposition method

in an increase in the computed bound. Putting aside the question of exactly how Step 3 is to be accomplished, the approach is straightforward in principle. Steps 1 and 2 are performed as in a traditional decomposition framework. Step 4 is accomplished by simply adding the newly generated inequalities to the list $[A'', b'']$ and reforming the appropriate bounding subproblem. Note that it is also possible to develop an analog based on an interpretation of the cutting plane method as a dynamic tightening of the implicit polyhedron [59]. In this case, the implicit polyhedron associated with the subproblem to be solved in Step 1 changes dynamically, rather than the list $[A'', b'']$. We have not yet investigated this class of methods.

Although some forms of this general method have appeared in the literature, they have received little attention thus far and naming conventions are not well-established. We would like to suggest here a naming convention that emphasizes the close relationship of these methods to each other. When the bounding subproblem is a Dantzig-Wolfe LP, we call the resulting method *price and cut*. When employed in a branch and bound framework, the overall technique is called *branch, price, and cut*. This method has been studied by a number of authors [12, 58, 36, 11, 57] and is described in more detail in Section 3.1. When the bounding subproblem is a Lagrangian dual, we call the method *relax and cut*. When relax and cut is used as the bounding procedure in a branch and bound framework, we call the overall method *branch, relax, and cut*. Variants of this method have been studied previously by several authors (see [43] for a survey) and is described in more detail in Section 3.2. Finally, in Section 3.3, we introduce a variant of the cutting plane method that employs a decomposition-based separation procedure. We call this method *decompose and cut* and

embed it within a branch and bound framework to obtain the method *branch, decompose, and cut*. Note that the distinction between what we call price and cut and what we call relax and cut is not that clear, since modern methods used for solving the subproblems in each case can be quite similar. The rough distinction we make between these methods, however, is the amount of primal solution information available. We will assume that when solving a Dantzig-Wolfe LP, the optimal DW decomposition $\hat{\lambda}$ is produced exactly. When solving a Lagrangian dual, only approximate primal solution information, if any at all, is typically available. We discuss solution methods for these subproblems in Section 3.4.

Step 3 is the crux of the method. In the context of the cutting plane method, this step is usually accomplished by applying one of the many known techniques for separating \hat{x}_{CP} from \mathcal{P} (see [1]). Violation of \hat{x}_{CP} is a necessary condition for an inequality to be improving, and hence such an inequality is *likely* to be effective. However, unless the inequality separates the entire optimal face F to the cutting plane LP, it will not be improving. Because we want to refer to these well-known results later in the paper, we state them formally as theorem and corollary without proof. See [56] for a thorough treatment of the theory of linear programming that leads to this result.

Theorem 2 *Let F be the face of optimal solutions to the cutting plane LP. Then $(a, \beta) \in \mathbb{R}^{n+1}$ is an improving inequality if and only if $a^\top y < \beta$ for all $y \in F$.*

Corollary 1 *If $(a, \beta) \in \mathbb{R}^{n+1}$ is an improving inequality and \hat{x} is an optimal solution to the cutting plane LP, then $a^\top \hat{x} < \beta$.*

Fortunately, even in the case when F is not separated in its entirety, the augmented cutting plane LP must have a different optimal solution, which in turn may be used to generate more potential improving inequalities. Since the condition of Theorem 2 is difficult to verify, one typically terminates the bounding procedure when increases resulting from additional inequalities become “too small.” In the next two sections, we examine how improving inequalities can be generated when the bounding subproblem is either a Dantzig-Wolfe LP or a Lagrangian dual. Then we return to the cutting plane method to discuss how decomposition can be used directly to improve our ability to solve the separation problem.

3.1 Price and Cut

Finding Improving Inequalities. Using the Dantzig-Wolfe LP as the bounding subproblem in the procedure of Figure 2 involves the simultaneous generation of columns¹ (pricing) and valid inequalities (cutting) during solution of the bounding subproblem. Simultaneous generation of both columns and valid inequalities is difficult in general because the addition of valid inequalities can destroy the structure of the column-generation subproblem (for a discussion of this, see [57]). Having solved the Dantzig-Wolfe LP, however, we can easily recover an optimal solution to the cutting plane LP using (8) and try to generate improving inequalities as in the cutting plane method. The generation of valid inequalities thus takes place in the original space and does not destroy the structure of the

¹To be more precise, we should use the term *variables* here. The term “column generation” implies a fixed number of rows.

column-generation subproblem in the Dantzig-Wolfe LP. This approach enables dynamic generation of valid inequalities while still retaining the bound improvement and other advantages yielded by Dantzig-Wolfe decomposition. A recent paper by Arão and Uchoa discusses this technique in more detail [21]. Because the same valid inequalities are generated with this method as would be generated in the cutting plane method, these two dynamic methods achieve the same bound. Despite the apparent symmetry, however, there is a fundamental difference between them. The optimal DW decomposition $\hat{\lambda}$ provides a decomposition of \hat{x}_{DW} into a convex combination of members of \mathcal{F}' . We refer to elements of \mathcal{F}' that have a positive weight in this combination as *members of the decomposition*. The following theorem shows how an optimal DW decomposition can be used to derive an alternate necessary condition for an inequality to be improving.

Theorem 3 *If $(a, \beta) \in \mathbb{R}^{(n+1)}$ is an improving inequality, then there must exist an $s \in \mathcal{F}'$ with $\hat{\lambda}_s > 0$ such that $a^\top s < \beta$.*

Proof. Let (a, β) be an improving inequality and suppose that $a^\top s \geq \beta$ for all $s \in \mathcal{F}'$ with $\hat{\lambda}_s > 0$. Since $\sum_{s \in \mathcal{F}'} \hat{\lambda}_s = 1$, we have $a^\top (\sum_{s \in \mathcal{F}'} s \hat{\lambda}_s) \geq \beta$. Hence, $a^\top \hat{x}_{DW} = a^\top (\sum_{s \in \mathcal{F}'} s \hat{\lambda}_s) \geq \beta$, which is a contradiction of Corollary 1, since \hat{x}_{DW} is an optimal solution to the cutting plane LP. ■

In other words, an inequality can be improving only if it is violated by at least one member of the decomposition. If \mathcal{I} is the set of all improving inequalities, then the following corollary is a restatement of Theorem 3.

Corollary 2 $\mathcal{I} \subseteq \mathcal{V} = \{(a, \beta) \in \mathbb{R}^{(n+1)} : a^\top s < \beta \text{ for some } s \in \mathcal{F}' \text{ such that } \hat{\lambda}_s > 0\}$

The importance of this result is that in many cases, it is easier to separate members of \mathcal{F}' from \mathcal{P} than to separate arbitrary real vectors. There are a number of well-known polyhedra for which the problem of separating an arbitrary real vector is difficult, but the problem of separating a solution to a given relaxation is easy. Some examples are discussed in Section 4. In Figure 3, we propose a new separation procedure to be embedded in price and cut that takes advantage of this fact.

The running time of this procedure depends in part on the cardinality of the decomposition. Carathéodory's Theorem assures us that there exists a decomposition with less than or equal to $\dim(\mathcal{P}) + 1$ members. Unfortunately, even if we limit our search to a particular known class of inequalities, the number of such inequalities violated by each member of \mathcal{D} in Step 2 may be extremely large and these inequalities may not be violated by \hat{x}_{DW} (such an inequality cannot be improving). Unless we enumerate *every* inequality in the set \mathcal{V} from Corollary 2, either implicitly or explicitly, the procedure does not guarantee that an improving inequality will be found, even if one exists. In cases where it is possible to examine the set \mathcal{V} in polynomial time, the worst-case complexity of the entire procedure is the same as that of optimizing over \mathcal{P}' . Obviously, it is unlikely that the set \mathcal{V} can be examined in polynomial time for situations in which the original separation problem is NP-complete. In such cases, the procedure to select inequalities that are likely to be violated by \hat{x}_{DW} in Step 2 is necessarily a problem-dependent heuristic. The effectiveness of this heuristic can be improved in a number of ways, some of which are discussed in [54]. Further details will be provided in a companion paper on the computational aspects of these methods.

Separation using a Dantzig-Wolfe Decomposition

Input: A DW decomposition $\hat{\lambda}$.

Output: A set $[D, d]$ of potentially improving inequalities for \mathcal{P} .

1. Form the set $\mathcal{D} = \{s \in \mathcal{F}' : \hat{\lambda}_s > 0\}$.
2. For each $s \in \mathcal{D}$, attempt to separate s from \mathcal{P} to obtain a set $[D^s, d^s]$ of violated inequalities.
3. Let $[D, d]$ be composed of the inequalities found in Step 2 that are also violated by \hat{x}_{DW} , so that $D\hat{x}_{DW} < d$.
4. Return $[D, d]$ as the set of potentially improving inequalities.

Figure 3: Separation using a Dantzig-Wolfe decomposition

Connections to Other Methods. By making connections to the cutting plane method, we can gain further insight. Consider the set

$$\mathcal{S} = \{s \in \mathcal{F}' : (c^\top - \hat{u}^\top A'')s = \hat{\alpha}\}, \quad (10)$$

where $\hat{\alpha}$ is again defined to be $z_{LR} - \hat{u}^\top b''$, so that $(\hat{u}, \hat{\alpha})$ is an optimal dual solution to the DWLP. Since \mathcal{S} is comprised exactly of those members of \mathcal{F}' corresponding to columns of the Dantzig-Wolfe LP with reduced cost zero, complementary slackness guarantees that the set \mathcal{S} must contain all members of the decomposition. The following theorem follows directly from this observation.

Theorem 4 *conv(S) is a face of \mathcal{P}' and contains \hat{x}_{DW} .*

Proof. We first show that *conv(S)* is a face of \mathcal{P}' . First, note that

$$(c^\top - \hat{u}^\top A'', \hat{\alpha})$$

defines a valid inequality for \mathcal{P}' since $\hat{\alpha}$ was defined as an optimal solution to the problem of minimizing over \mathcal{P}' with objective function $c^\top - \hat{u}^\top A''$. Thus, the set

$$G = \{s \in \mathcal{P}' : (c^\top - \hat{u}^\top A'')s = \hat{\alpha}\}, \quad (11)$$

is a face of \mathcal{P}' that contains \mathcal{S} . We claim that *conv(S)* = G . Since G is convex and contains \mathcal{S} , it also contains *conv(S)*, so we just need to show that *conv(S)* contains G . We do so by showing that the extreme points of G are among those of \mathcal{S} . By construction, all extreme points of \mathcal{P}' are members of \mathcal{F}' and the extreme points of G are also extreme points of \mathcal{P}' . Therefore, the extreme points of G must be members of \mathcal{F}' and contained in \mathcal{S} . The claim follows and *conv(S)* is a face of \mathcal{P}' .

The fact that $\hat{x}_{DW} \in \text{conv}(\mathcal{S})$ follows from the fact that \hat{x}_{DW} is a convex combination of members of \mathcal{S} . ■

An important consequence of this result is contained in the following corollary.

Corollary 3 *If F is the face of optimal solutions to the cutting plane LP, then $F \subseteq \text{conv}(\mathcal{S}) \cap \mathcal{Q}''$.*

Proof. Any optimal solution \hat{x}_{CP} to the cutting plane LP must be contained in $\mathcal{P}' \cap \mathcal{Q}''$ by definition and must also satisfy

$$c^\top \hat{x}_{CP} = z_{CP} = z_{LD} = \hat{\alpha} + \hat{u}^\top b'' = \hat{\alpha} + \hat{u}^\top A'' \hat{x}_{CP}, \quad (12)$$

where the last equality in this chain is a consequence of complementary slackness. ■

Hence, the convex hull of the decomposition is a subset of $\text{conv}(\mathcal{S})$ that contains \hat{x}_{DW} and can be thought of as a surrogate for the face of optimal solutions to the cutting plane LP. Combining this corollary with Theorem 2, we conclude that separation of \mathcal{S} from \mathcal{P} is a sufficient condition for an inequality to be improving. Although this sufficient condition is difficult to verify in practice, it does provide additional motivation for the method described in Figure 3.

The convex hull of \mathcal{S} is typically a proper face of \mathcal{P}' . It is possible, however, for \hat{x}_{DW} to be an inner point of \mathcal{P}' .

Theorem 5 *If \hat{x}_{DW} is an inner point of \mathcal{P}' , then $\text{conv}(\mathcal{S}) = \mathcal{P}'$.*

Proof. We prove the contrapositive. Suppose $\text{conv}(\mathcal{S})$ is a proper face of \mathcal{P}' . Then there exists a facet-defining valid inequality $(a, \beta) \in \mathbb{R}^{n+1}$ such that $\text{conv}(\mathcal{S}) \subseteq \{x \in \mathbb{R}^n : ax = \beta\}$. By Theorem 4, $\hat{x}_{DW} \in \text{conv}(\mathcal{S})$ and \hat{x}_{DW} therefore cannot satisfy the definition of an inner point. ■

In this case, illustrated graphically in Figure 4(a), $z_{DW} = z_{LP}$ and Dantzig-Wolfe decomposition does not improve the bound. All columns of the Dantzig-Wolfe LP have reduced cost zero and any extremal member of \mathcal{F}' can be made a member of the decomposition. A necessary condition for an optimal fractional solution to be an inner point of \mathcal{P}' is that the dual value of the convexity constraint in an optimal solution to the Dantzig-Wolfe LP be zero. This condition indicates that the chosen relaxation may be too weak.

A second case of potential interest is when $F = \text{conv}(\mathcal{S}) \cap \mathcal{Q}''$, illustrated graphically in Figure 4(b). This condition can be detected by examining the objective function values of the members of the decomposition. If they are all identical, any member of the decomposition that is contained in \mathcal{Q}'' (if one exists) must be optimal for the original ILP, since it is feasible and has objective function value equal to z_{LP} . In this case, all constraints of the Dantzig-Wolfe LP *other than* the convexity constraint must have dual value zero, since removing them does not change the optimal solution value. The more typical case, in which F is a proper subset of $\text{conv}(\mathcal{S}) \cap \mathcal{Q}''$, is shown in Figure 4(c).

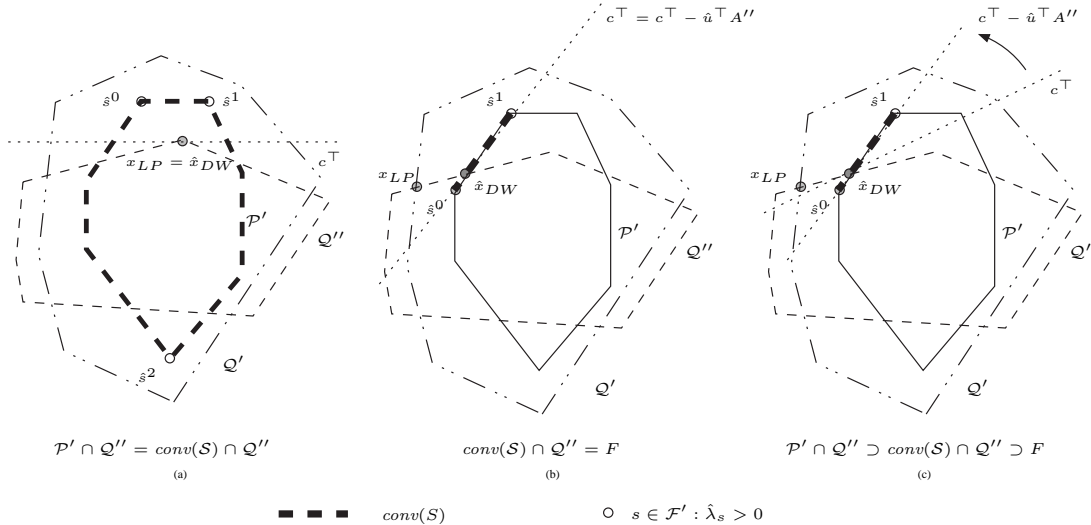


Figure 4: The relationship of $\mathcal{P}' \cap \mathcal{Q}''$, $\text{conv}(\mathcal{S})$, and the face F , for different cost vectors.

3.2 Relax and Cut

Finding Improving Inequalities. When the bounding subproblem is the Lagrangian dual, it is more difficult to obtain the primal solution information readily available to us with both the cutting plane method and price and cut. The amount of primal solution information depends on the algorithm used to solve the Lagrangian dual. With methods such as the volume algorithm [9], it is possible to obtain an approximate primal solution. For the sake of discussion, however, we assume in what follows that no primal solution information is available. In such a case, we can attempt to separate $\hat{s} = \text{argmin } z_{LR}(\hat{u})$ from \mathcal{P} , where \hat{u} is the vector of optimal dual multipliers defined earlier. Since \hat{s} is a member of \mathcal{F}' , we are again taking advantage of our ability to separate members of \mathcal{F}' from \mathcal{P} effectively. If successful, we immediately “dualize” this new constraint by adding it to $[A'', b'']$, as described in Section 3. This has the effect of introducing a new dual multiplier and slightly perturbing the Lagrangian objective function.

As with both the cutting plane and price and cut methods, the difficulty with relax and cut is that the valid inequalities generated by separating \hat{s} from \mathcal{P} may not be improving, as Guignard first observed in [31]. Furthermore, we cannot verify the condition of Corollary 1, which is the best available necessary condition for an inequality to be improving. To deepen understanding of the potential effectiveness of the valid inequalities generated during relax and cut, we further examine the relationship between \hat{s} and \hat{x}_{DW} .

Connections to Other Methods. By considering again the reformulation of the Lagrangian dual as the linear program (6), we observe that each constraint binding at an optimal solution corresponds to an alternative optimal solution to the Lagrangian subproblem with multipliers \hat{u} . The binding constraints of (6) correspond to variables with reduced cost zero in the Dantzig-Wolfe LP (7), so it follows immediately that the set \mathcal{S} from (10) is

also the set of all alternative solutions to $z_{LR}(\hat{u})$.

Because \hat{x}_{DW} is both an optimal solution to the cutting plane LP and is contained in $\text{conv}(\mathcal{S})$, it also follows that

$$c^\top \hat{x}_{DW} = c^\top \hat{x}_{DW} + \hat{u}^\top (b'' - A'' \hat{x}_{DW}) = (c^\top - \hat{u}^\top A'') \hat{x}_{DW} + \hat{u}^\top b''.$$

In other words, the penalty term in the objective function of the Lagrangian subproblem (4) serves to rotate the original objective function so that it becomes parallel to the face $\text{conv}(\mathcal{S})$, while the constant term $\hat{u}^\top b''$ ensures that \hat{x}_{DW} has the same cost with both the original and the Lagrangian objective function. This is illustrated in Figure 4(c).

One conclusion that can be drawn from these observations is that solving the Dantzig-Wolfe LP produces a *set* of alternative optimal solutions to the Lagrangian subproblem with multipliers \hat{u} , at least one of which must be violated by a given improving inequality. This yields a verifiable necessary condition for a generated inequality to be improving. Relax and cut, in its most straightforward incarnation, produces one member of this set. In this case, even if improving inequalities exist, it is possible that none of them are violated by the member of \mathcal{S} so produced, especially if it has a small weight in the optimal DW decomposition. Note, however, that by keeping track of the solutions to the Lagrangian subproblem that are produced while solving the Lagrangian dual, one can approximate the optimal DW decomposition produced by solving the Dantzig-Wolfe LP. This is the approach taken by the volume algorithm [9] and other subgradient-based methods. As in price and cut, when \hat{x}_{DW} is an inner point of \mathcal{P}' , the decomposition does not improve the bound and all members of \mathcal{F}' are alternative optimal solutions to the Lagrangian subproblem. This situation is depicted in Figure 4(a). In this case, separating an optimal solution to $z_{LR}(\hat{u})$ from \mathcal{P} is unlikely to yield an improving inequality.

3.3 Decompose and Cut

The use of an optimal DW decomposition to aid in separation is easy to extend to a traditional branch and cut framework using a technique we call *decompose and cut*, originally proposed in [53] and further developed in [37] and [54]. Suppose now that we are given an optimal fractional solution \hat{x}_{CP} obtained directly by solving the cutting plane LP and suppose that given $s \in \mathcal{F}'$, we can determine effectively whether $s \in \mathcal{F}$ and if not, generate a valid inequality (a, β) violated by s . By first *decomposing* \hat{x}_{CP} (i.e., expressing \hat{x}_{CP} as a convex combination of members of \mathcal{F}') and then separating each member of this decomposition from \mathcal{P} in a fashion analogous to that described in Figure 3, we can separate \hat{x}_{CP} from the polyhedron \mathcal{P} .

The difficult step is finding the decomposition of \hat{x}_{CP} . This can be accomplished by solving a linear program whose columns are the members of \mathcal{F}' , as described in Figure 5. This linear program is reminiscent of a Dantzig-Wolfe LP and in fact can be solved using an analogous column-generation scheme, as described in Figure 6. This scheme can be seen as the “inverse” of the method described in Section 3.1, since it begins with the fractional solution \hat{x}_{CP} and tries to compute a decomposition, instead of the other way around. By the equivalence of optimization and separation, we conclude that the worst-case complexity of finding a decomposition of \hat{x}_{CP} is the same as that of optimizing over \mathcal{P}' .

Separation in Decompose and Cut

Input: $\hat{x} \in \mathbb{R}^n$

Output: A valid inequality for \mathcal{P} violated by \hat{x} , if one is found.

1. Apply standard separation techniques to separate \hat{x} . If one of these returns a violated inequality, then STOP and output the violated inequality.
2. Otherwise, solve the linear program

$$\max_{\lambda \in \mathbb{R}_+^{\mathcal{F}'}} \{ \mathbf{0}^\top \lambda : \sum_{s \in \mathcal{F}'} s \lambda_s = \hat{x}, \sum_{s \in \mathcal{F}'} \lambda_s = 1 \}, \quad (13)$$

as in Figure 6

3. The result of Step 2 is either (1) a subset \mathcal{D} of members of \mathcal{F}' participating in a convex combination of \hat{x} , or (2) a valid inequality (a, β) for \mathcal{P} that is violated by \hat{x} . In the first case, go to Step 4. In the second case, STOP and output the violated inequality.
4. Attempt to separate each member of \mathcal{D} from \mathcal{P} . For each inequality violated by a member of \mathcal{D} , check whether it is also violated by \hat{x} . If an inequality violated by \hat{x} is encountered, STOP and output it.

Figure 5: Separation in the decompose and cut method

Column Generation in Decompose and Cut

Input: $\hat{x} \in \mathbb{R}^n$

Output: Either (1) a valid inequality for \mathcal{P} violated by \hat{x} ; or (2) a subset \mathcal{D} of \mathcal{F}' and a vector $\hat{\lambda} \in \mathbb{R}_+^{\mathcal{D}}$ such that $\sum_{s \in \mathcal{D}} \lambda_s s = \hat{x}$ and $\sum_{s \in \mathcal{D}} \lambda_s = 1$.

- 2.0 Generate an initial subset \mathcal{G} of \mathcal{F}' .
- 2.1 Solve (13), replacing \mathcal{F}' by \mathcal{G} . If this linear program is feasible, then the elements of \mathcal{F}' corresponding to the nonzero components of $\hat{\lambda}$, the current solution, comprise the set \mathcal{D} , so STOP.
- 2.2 Otherwise, let (a, β) be a valid inequality for $\text{conv}(\mathcal{G})$ violated by \hat{x} . Solve the optimization problem over \mathcal{P}' with cost vector a and let s be the resulting solution. If the optimal value is less than β , then add s to \mathcal{G} and go to 2.1. Otherwise, (a, β) is an inequality valid for $\mathcal{P}' \supseteq \mathcal{P}$ and violated by \hat{x} , so STOP.

Figure 6: Column generation for the decompose and cut method

Once the decomposition is found, it can be used as before to locate a violated inequality using the methodology discussed earlier. In contrast to price and cut, it is possible that $\hat{x}_{CP} \notin \mathcal{P}'$. This could occur, for instance, if exact separation methods for \mathcal{P}' are too expensive to apply consistently. In this case, it is obviously not possible to find a decomposition in Step 2 of Figure 5. The proof of infeasibility for the linear program (13), however, provides an inequality separating \hat{x}_{CP} from \mathcal{P}' at no additional expense. Hence, even if we fail to find a decomposition, we still find an inequality valid for \mathcal{P} and violated by \hat{x}_{CP} . This idea was originally suggested in [53] and was further developed by Kopman in [37]. A similar concept was also discovered and developed independently by Applegate, et al. [3].

Applying decompose and cut in every iteration as the sole means of separation is theoretically equivalent to price and cut. In practice, however, the decomposition is only computed when needed, i.e., when less expensive separation heuristics fail to separate the optimal fractional solution. This could give decompose and cut an edge in terms of computational efficiency. In other respects, the computations performed in each method are similar.

3.4 Implementation and Extensions

In practice, there are many variations on the general theme described here. The details surrounding implementation of these methods will be covered in a separate paper, but we would like to give some taste for the issues involved. An important aspect of the implementation of these methods is the algorithm used for solving the subproblem P^i in Step 2 of the algorithm in Figure 2. Three general categories of methods for solving such subproblems are simplex methods, interior point methods, and subgradient methods. Simplex methods provide accurate primal solution information, but updates to the dual solution each iteration are relatively expensive. In their most straightforward form, they also tend to converge slowly when used to solve a Dantzig-Wolfe LP by column generation because of the fact that they produce basic dual solutions, which can change substantially from one iteration to the next. This problem can be addressed by implementing one of a number of stabilization methods that prevent the dual solution from changing “too much” from one iteration to the next (for a survey, see [40]). In their most straightforward form, subgradient methods do not produce primal solution information. However, it is possible to extract approximate primal solutions from variants of subgradient such as the volume algorithm [9]. Subgradient methods also have convergence issues without some form of stabilization. A recent class of algorithms that has proven effective in this regard is bundle methods [18]. Interior point methods may provide a middle ground by providing accurate primal solution information and more stable dual solutions [55, 28]. In addition, hybrid methods that alternate between simplex and subgradient methods for updating the dual solution have also shown promise [10, 33].

Inner methods for solving the bounding subproblem can be viewed essentially as alternating between updating the dual solution and using that solution to produce further information about \mathcal{P}' by solving a “pricing” subproblem. Further variants can be generated by not insisting on a “complete” dual update before solving the pricing problem [29, 17]. Also of particular interest are classes of methods obtained by not solving P^i to optimality

before generating cuts [43]. The software framework we are developing will allow essentially any sequence of dual updates, pricing, and cutting by any of the methods discussed here. This leads to a wide range of possibilities, very few of which have been investigated in the literature so far.

4 Applications

In this section, we illustrate the concepts presented so far with three examples. For each example, we discuss the key elements needed to apply the framework: (1) the original ILP formulation, (2) the explicit and implicit polyhedra, and (3) known classes of valid inequalities that can be dynamically generated in the fashion we have described. The concept of a “class” of inequalities comes from the well-known *template paradigm* for separation, so named by Applegate et al. [3]. This paradigm operates on the precept that it may be possible to solve the separation problem for a given class of inequalities valid for the polyhedron \mathcal{P} effectively, though the general separation problem for \mathcal{P} is difficult. Our framework extends this paradigm by considering classes of valid inequalities along with relaxations for which the separation of an arbitrary real vector is difficult but separation of solutions to the relaxation can be accomplished effectively. In addition to the three examples presented here, we have discovered a number of common ILPs with classes of valid inequalities and relaxations that fit into this framework, such as the Generalized Assignment Problem [51], the Edge-Weighted Clique Problem [34], the Traveling Salesman Problem [5], the Knapsack Constrained Circuit Problem [39], the Rectangular Partition Problem [16], the Linear Ordering Problem [15], and the Capacitated Minimum Spanning Tree Problem [24].

4.1 Vehicle Routing Problem

We first consider the *Vehicle Routing Problem* (VRP) introduced by Dantzig and Ramser [20]. In this \mathcal{NP} -hard optimization problem, a fleet of k vehicles with uniform capacity C must service known customer demands for a single commodity from a common depot at minimum cost. Let $V = \{1, \dots, n\}$ index the set of customers and let the depot have index 0. Associated with each customer $i \in V$ is a demand d_i . The cost of travel from customer i to j is denoted c_{ij} and we assume that $c_{ij} = c_{ji} > 0$ if $i \neq j$ and $c_{ii} = 0$.

By constructing an associated complete undirected graph G with vertex set $N = V \cup \{0\}$ and edge set E , we can formulate the VRP as an integer program. A *route* is a set of nodes $R = \{i_1, i_2, \dots, i_m\}$ such that the members of R are distinct. The edge set of R is $E_R = \{\{i_j, i_{j+1}\} : j \in 0, \dots, m\}$, where $i_0 = i_{m+1} = 0$. A feasible solution is then any subset of E that is the union of the edge sets of k disjoint routes $R_i, i \in [1..k]$, each satisfying the capacity restriction, i.e., $\sum_{j \in R_i} d_j \leq C, \forall i \in [1..k]$. Each route corresponds to a set of customers serviced by one of the k vehicles. To simplify the presentation, let us define some additional notation. Let $\delta(S) = \{\{i, j\} \in E : i \in S, j \notin S\}$, $E(S : T) = \{\{i, j\} : i \in S, j \in T\}$, $E(S) = E(S : S)$ and $x(F) = \sum_{e \in F} x_e$.

By associating a variable with each edge in the graph, we obtain the following formula-

tion of this ILP [38]:

$$\min \sum_{e \in E} c_e x_e$$

$$x(\delta(\{0\})) = 2k \tag{14}$$

$$x(\delta(\{v\})) = 2 \quad \forall v \in V, \tag{15}$$

$$x(\delta(S)) \geq 2b(S) \quad \forall S \subseteq V, |S| > 1, \tag{16}$$

$$x_e \in \{0, 1\} \quad \forall e \in E(V), \tag{17}$$

$$x_e \in \{0, 1, 2\} \quad \forall e \in \delta(0). \tag{18}$$

Here, $b(S)$ represents a lower bound on the number of vehicles required to service the set of customers S . Inequalities (14) ensure that there are exactly k vehicles, each departing from and returning to the depot, while inequalities (15) require that each customer must be serviced by exactly one vehicle. Inequalities (16), known as the *generalized subtour elimination constraints* (GSECs) can be viewed as a generalization of the subtour elimination constraints from the *Traveling Salesman Problem* (TSP) and enforce connectivity of the solution, as well as ensuring that no route has total demand exceeding capacity C . For ease of computation, we can define $b(S) = \lceil (\sum_{i \in S} d_i) / C \rceil$, a trivial lower bound on the number of vehicles required to service the set of customers S .

Returning to our earlier notation and setup, the set of feasible solutions to the VRP is

$$\mathcal{F} = \{x \in \mathbb{R}^E : x \text{ satisfies (14) – (18)}\}$$

and we call $\mathcal{P} = \text{conv}(\mathcal{F})$ the *VRP polytope*. Many classes of valid inequalities for the VRP polytope have been reported in the literature (see [48] for a survey). Significant effort has been devoted to developing efficient algorithms for separating an arbitrary fractional point using these classes of inequalities (see [35] for recent results).

We will concentrate here on the separation of GSECs. The separation problem for GSECs was shown to be in \mathcal{NP} -complete by Harche and Rinaldi (see [4]), even when $b(S)$ is taken to be $\lceil (\sum_{i \in S} d_i) / C \rceil$. In [35], Lysgaard, et al. review heuristic procedures for separating arbitrary fractional solutions from the GSEC polyhedron. Although GSECs are part of the formulation presented above, there are exponentially many of them, so they must either be generated dynamically or included as part of the description of the implicit polyhedron. We discuss three alternatives for the implicit polyhedron: the *Multiple Traveling Salesman* polytope, the *Perfect b-Matching* polytope, and the *Degree-constrained k-Tree* polytope. For each of these alternatives, solutions to the relaxation can be easily separated from the GSEC polyhedron.

Perfect b-Matching Problem. With respect to the graph G , the *Perfect b-Matching Problem* is to find a minimum weight subgraph of G such that $x(\delta(v)) = b_v \quad \forall v \in V$. This

problem can be formulated by dropping the GSECs from the VRP formulation, resulting in the implicit polyhedron as $\mathcal{P}' = \text{conv}(\mathcal{F}')$, where

$$\mathcal{F}' = \{x \in \mathbb{R}^E : x \text{ satisfies (14), (15), (17), (18)}\}.$$

In [47], Müller-Hannemann and Schwartz, present several fast polynomial algorithms for solving b -Matching. Note that the explicit polyhedron is \mathbb{R}^n in this case.

In [46], Miller uses the b -Matching relaxation to solve the VRP by branch, relax, and cut. He suggests separating b -matchings from the GSEC polytope as follows. Consider a member s of \mathcal{F}' and its support graph G_s (a b -Matching). If G_s is disconnected, then each component immediately induces a violated GSEC. On the other hand, if G_s is connected, we first remove the edges incident to the depot vertex and find the connected components, which comprise the routes described earlier. To identify a violated GSEC, we compute the total demand of each route, checking whether it exceeds capacity. If not, the solution is feasible for the original ILP and does not violate any GSECs. If so, the set S of customers on any route whose total demand exceeds capacity induces a violated GSEC. This separation routine runs in $O(n)$ time and can be used in any of the integrated decomposition methods previously described. Figure 7 gives an optimal fractional solution (a) to an LP relaxation of the VRP expressed as a convex combination of two b -Matchings (b,c). In this example, the capacity $C = 35$ and by inspection we find a violated GSEC in the second b -Matching (c) with S equal to the marked component. This inequality is also violated by the optimal fractional solution, since $\hat{x}(\delta(S)) = 3.0 < 4.0 = 2b(S)$.

Minimum Degree-constrained k -Tree Problem. A k -tree is defined as a spanning subgraph of G that has $n + k$ edges (recall that G has $n + 1$ nodes). A *degree-constrained k -tree* (k -DCT), as defined by Fisher in [23], is a k -tree with degree $2k$ at node 0. The *Minimum k -DCT Problem* is that of finding a minimum cost k -DCT, where the cost of a k -DCT is the sum of the costs on the edges present in the k -DCT. Fisher [23] introduced this relaxation as part of a Lagrangian relaxation-based algorithm for solving the VRP.

The k -DCT polyhedron is obtained by first adding the redundant constraint

$$x(E) = n + k, \tag{19}$$

then deleting the degree constraints (15), and finally, relaxing the capacity to $C = \sum_{i \in S} d_i$. Relaxing the capacity constraints gives $b(S) = 1$ for all $S \subseteq V$, and replaces the set of constraints (16) with

$$\sum_{e \in \delta(S)} x_e \geq 2, \forall S \subseteq V, |S| > 1. \tag{20}$$

The implicit polyhedron is then defined as $\mathcal{P}' = \text{conv}(\mathcal{F}')$, where

$$\mathcal{F}' = \{x \in \mathbb{R}^E : x \text{ satisfies (14), (16), (17), (19)}\}.$$

The explicit polyhedron is described simply by the constraints (15). In [60], Wei and Yu give a polynomial algorithm for solving the Minimum k -DCT Problem that runs in $O(n^2 \log n)$ time. In [45], Martinhon et al. study the use of the k -DCT relaxation for the VRP in

the context branch, relax and cut. Again, consider separating a member s of \mathcal{F}' from the polyhedron defined by all GSECS. It is easy to see that for GSECS, an algorithm identical to that described above can be applied. Figure 7 also gives the optimal fractional solution (a) expressed as a convex combination of four k -DCTs (d-g). Removing the depot edges, and checking each components demand, we easily identify the violated GSEC in k -DCT (g).

Multiple Traveling Salesman Problem. The *Multiple Traveling Salesman Problem* (k -TSP) is an uncapacitated version of the VRP obtained by adding the degree constraints to the k -DCT polyhedron. The implicit polyhedron is then defined as $\mathcal{P}' = \text{conv}(\mathcal{F}')$, where

$$\mathcal{F}' = \{x \in \mathbb{R}^E : x \text{ satisfies (14), (15), (17), (18), (20)}\}.$$

Although the k -TSP is an \mathcal{NP} -hard optimization problem, small instances can be solved effectively by transformation into an equivalent TSP obtained by adjoining to the graph $k-1$ additional copies of node 0 and its incident edges. In this case, the explicit polyhedron is again \mathbb{R}^n . In [54], Ralphs et al. report on an implementation of branch, decompose and cut using the k -TSP as a relaxation.

4.2 Axial Assignment Problem

The *Axial Assignment Problem* (AAP) is that of finding a minimum-weight clique cover of the complete tri-partite graph $K_{n,n,n}$. Let I, J and K be three disjoint sets with $|I| = |J| = |K| = n$ and set $H = I \times J \times K$. AAP can be formulated as the following binary integer program:

$$\begin{aligned} \min \quad & \sum_{(i,j,k) \in H} c_{ijk} x_{ijk} \\ & \sum_{(j,k) \in J \times K} x_{ijk} = 1 \quad \forall i \in I \end{aligned} \tag{21}$$

$$\sum_{(i,k) \in I \times K} x_{ijk} = 1 \quad \forall j \in J \tag{22}$$

$$\sum_{(i,j) \in I \times J} x_{ijk} = 1 \quad \forall k \in K \tag{23}$$

$$x_{ijk} \in \{0, 1\} \quad \forall (i, j, k) \in H \tag{24}$$

A number of applications of AAP can be found in the literature (see Piersjalla [18,19]). AAP is known to be \mathcal{NP} -hard [25]. As before, the set of feasible solutions to AAP is noted as

$$\mathcal{F} = \{x \in \mathbb{R}^H : x \text{ satisfies (21) - (24)}\}$$

and we set $\mathcal{P} = \text{conv}(\mathcal{F})$.

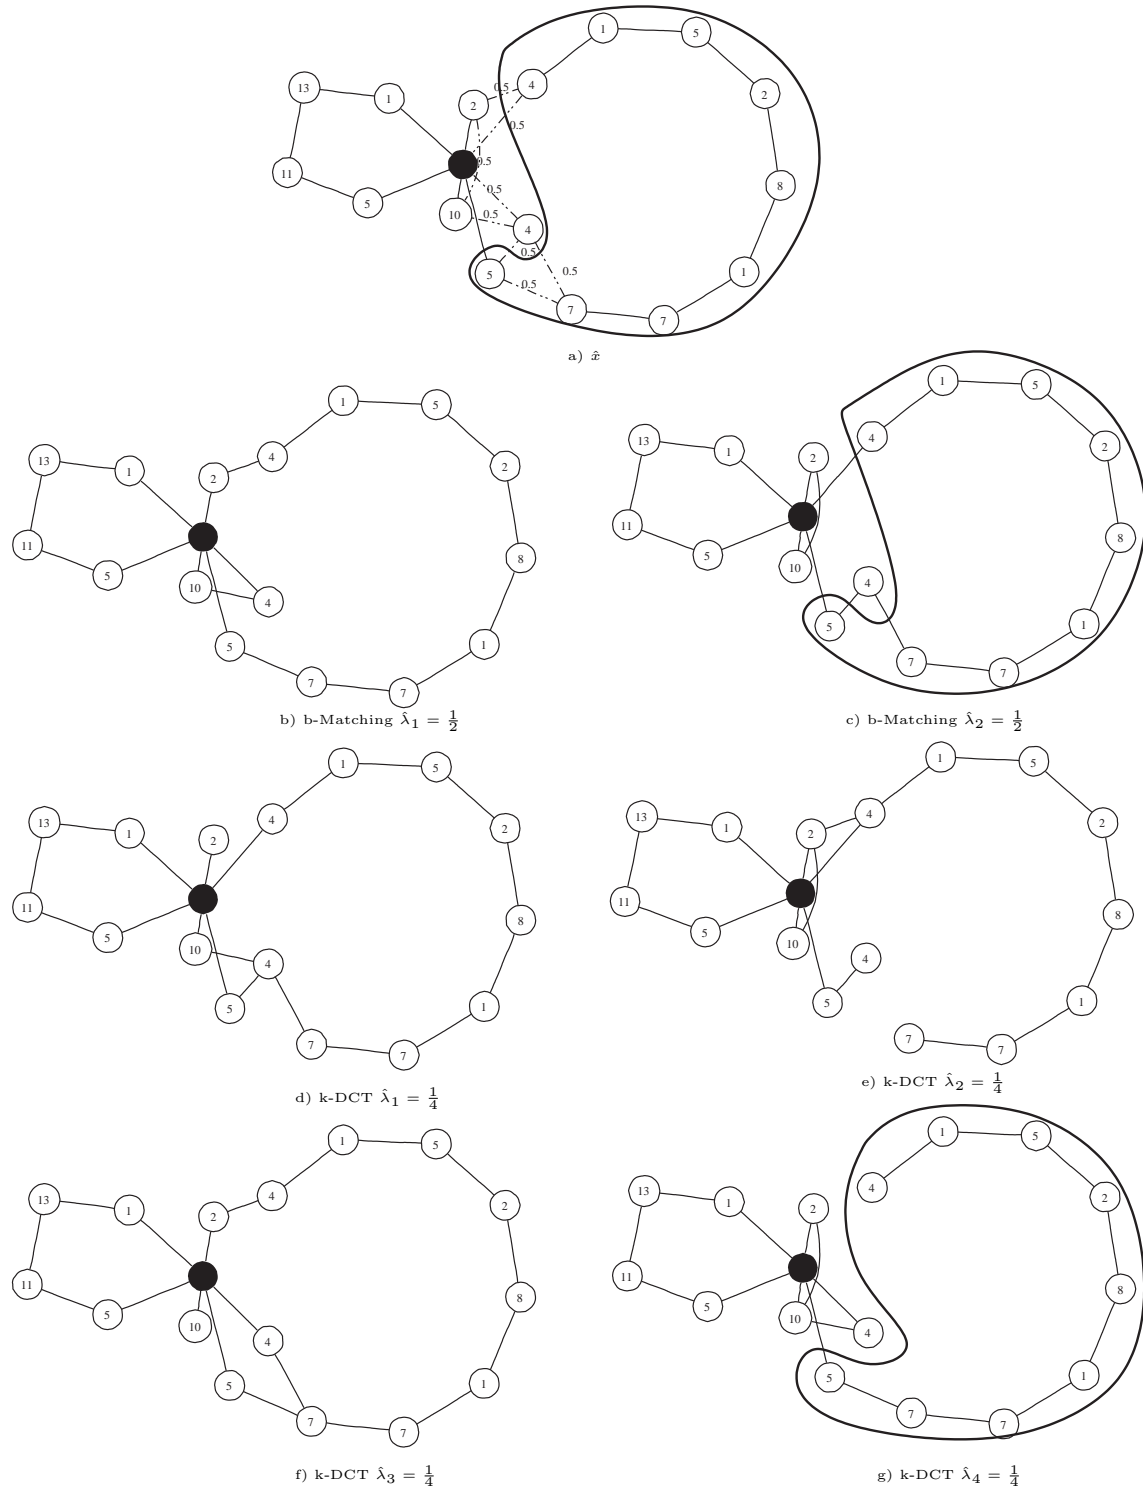


Figure 7: Example of a decomposition into *b*-Matchings and *k*-DCTs




Figure 8: Example of a decomposition into assignments.

In [7], Balas and Saltzman study the polyhedral structure of \mathcal{P} and introduce several classes of facet-inducing inequalities. Let $u, v \in H$ and define $|u \cap v|$ to be the numbers of coordinates for which the vectors u and v have the same value. Let $C(u) = \{w \in H : |u \cap w| = 2\}$ and $C(u, v) = \{w \in H : |u \cap w| = 1, |w \cap v| = 2\}$. We consider two classes of facet-inducing inequalities $Q_1(u)$ and $P_1(u, v)$ for \mathcal{P} ,

$$x_u + \sum_{w \in C(u)} x_w \leq 1 \quad \forall u \in H, \quad (25)$$

$$x_u + \sum_{w \in C(u, v)} x_w \leq 1 \quad \forall u, v \in H, |u \cap v| = 0. \quad (26)$$

Note that these include all the clique facets of the intersection graph of $K_{n,n,n}$ [7]. In [6], Balas and Qi describe algorithms that solve the separation problem for the polyhedra defined by the inequalities $Q_1(u)$ and $P_1(u, v)$ in $O(n^3)$ time.

Balas and Saltzman consider the use of the classical Assignment Problem (AP) as a relaxation of AAP in an early implementation of branch, relax, and cut [8]. Following their lead, we define the implicit polyhedron to be $\mathcal{P}' = \text{conv}(\mathcal{F}')$, where

$$\mathcal{F}' = \{x \in \mathbb{R}^H : x \text{ satisfies (22) - (24)}\}.$$

The explicit polyhedron is then described by constraints (21). The AP can be solved in $O(n^{5/2} \log(nC))$ time where $C = \max_{w \in H} c_w$, by the cost-scaling algorithm [2]. Consider

separating a member s of \mathcal{F}' from the polyhedron defined by $Q_1(u)$ for all $u \in H$. Let $L(s)$ be the set of n triplets corresponding to the nonzero components of s (the assignment from J to K). It is easy to see that if there exist $u, v \in L(s)$ such that $u = (i_0, j_0, k_0)$ and $v = (i_0, j_1, k_1)$, i.e., the assignment *overcovers* the set I , then both $Q(i_0, j_0, k_1)$ and $Q(i_0, j_1, k_0)$ are violated by s . Figure 8 shows the decomposition of an optimal fractional solution \hat{x} (a) into a convex combination of assignments (b-d). In the second assignment (c), the pair of triplets $(0, 3, 1)$ and $(0, 0, 3)$ satisfies the condition just discussed and identifies two violated valid inequalities, $Q_1(0, 3, 3)$ and $Q_1(0, 0, 1)$ which violate \hat{s} . The latter also violates \hat{x} and is illustrated in (e). This separation routine runs in $O(n)$ time.

Now consider separating a member s of \mathcal{F}' from the polyhedron defined by $P_1(u, v)$ for all $u, v \in H$. As above, for any pair of assignments that correspond to nonzero components of s and have the form $(i_0, j_0, k_0), (i_0, j_1, k_1)$, we know s violates $P_1((i_0, j_0, k_0), (i_0, j_1, k_1))$, $\forall i \neq i_0$ and $P_1((i_0, j_1, k_1), (i_0, j_0, k_0))$, $\forall i \neq i_0$. In Figure 8, the second assignment (c) is violated by $P_1((0, 0, 3), (1, 3, 1))$. This inequality is also violated by \hat{x} and is illustrated in (f). Once again, this separation routine runs in $O(n)$ time.

4.3 Steiner Tree Problem

Let $G = (V, E)$ be a complete undirected graph with vertex set $V = \{1, \dots, n\}$, edge set E and a positive weight c_e associated with each edge $e \in E$. Let $T \subset V$ define the set of *terminals*. The *Steiner Tree Problem* (STP), which is \mathcal{NP} -hard, is that of finding a subgraph that spans T (called a *Steiner tree*) and has minimum edge cost. In [13], Beasley formulated the STP as a side constrained *Minimum Spanning Tree Problem* (MSTP) as follows. Let $r \in T$ be some terminal and define an artificial vertex 0. Now, construct the augmented graph $\bar{G} = (\bar{V}, \bar{E})$ where $\bar{V} = V \cup \{0\}$ and $\bar{E} = E \cup \{\{i, 0\} : i \in (V \setminus T) \cup \{r\}\}$. Let $c_{i0} = 0$ for all $i \in (V \setminus T) \cup \{r\}$. Then, the STP is equivalent to finding a minimum spanning tree (MST) in \bar{G} subject to the additional restriction that any vertex $i \in (V \setminus T)$ connected by edge $\{i, 0\} \in \bar{E}$ must have degree one.

By associating a binary variable x_e with each edge $e \in \bar{E}$, indicating whether or not the edge is selected, we can then formulate the STP as the following integer program:

$$\min \sum_{e \in E} c_e x_e$$

$$x(\bar{E}) = |\bar{V}| - 1 \tag{27}$$

$$x(E(S)) \leq |S| - 1 \quad \forall S \subseteq \bar{V} \tag{28}$$

$$x_{i0} + x_e \leq 1 \quad \forall e \in \delta(i), i \in (V \setminus T) \tag{29}$$

$$x_e \in \{0, 1\} \quad \forall e \in \bar{E} \tag{30}$$

Inequalities (27) and (28) ensure that the solution forms a spanning tree on \bar{G} . Inequalities (28) are subtour elimination constraints (similar to those used in the TSP). Inequalities (29) are the side constraints that ensure the solution can be converted to a Steiner tree by dropping the edges in $\bar{E} \setminus E$.

The set of feasible solutions to the STP is

$$\mathcal{F} = \{x \in \mathbb{R}^{\bar{E}} : x \text{ satisfies (27) -- (30)}\}.$$

We set $\mathcal{P} = \text{conv}(\mathcal{F})$ as before. We consider two classes of valid inequalities for \mathcal{P} that are lifted versions of the subtour elimination constraints (SEC).

$$x(E(S)) + x(E(S \setminus T : \{0\})) \leq |S| - 1 \quad \forall S \subseteq V, S \cap T \neq \emptyset, \quad (31)$$

$$x(E(S)) + x(E(S \setminus \{v\} : \{0\})) \leq |S| - 1 \quad \forall S \subseteq V, S \cap T = \emptyset, v \in S \quad (32)$$

The class of valid inequalities (31) were independently introduced by Goemans [27], Lucena [41] and Margot, Prodon, and Liebling [44], for another extended formulation of STP. The inequalities (32) were introduced in [27, 44]. The separation problem for inequalities (31) and (32) can be solved in $O(n^4)$ time through a series of max-flow computations.

In [42], Lucena considers the use of MSTP as a relaxation of STP in the context of a branch, relax, and cut algorithm. Once again, we cannot include inequalities (28) in the description of the explicit polyhedron, which is described by constraints (27) and (29). The implicit polyhedron is defined to be $\mathcal{P}' = \text{conv}(\mathcal{F}')$, where

$$\mathcal{F}' = \{x \in \mathbb{R}^{\bar{E}} : x \text{ satisfies (27), (28), (30)}\}.$$

The MSTP can be solved in $O(n^2 \log n)$ time using Prim's algorithm [52]. Consider the separation of a member of $s \in \mathcal{F}'$ from the polyhedron defined by the lifted subtour inequalities (31) and (32). In order to identify a violated inequality of the form (31) or (32) we remove the artificial vertex 0 and find the connected components on the resulting subgraph. Any component of size greater than 1 which does not contain r and does contain a terminal, defines a violated SEC (31). In addition, if the component does not contain any terminals, then each vertex in the component which was not connected to the artificial vertex defines a violated SEC (32).

Figure 9 gives an optimal fractional solution \hat{x} (a) to an LP relaxation of the STP expressed as a convex combination of two MSTs (b,c). In this figure, the artificial vertex is black, the terminals are gray and $r = 3$. By removing the artificial vertex, we easily find a violated SEC in the second spanning tree (c) with S equal to the marked component. This inequality is also violated by the optimal fractional solution, since $\hat{x}(E(S)) + \hat{x}(E(S \setminus T : \{0\})) = 3.5 > 3 = |S| - 1$. It should also be noted that the first spanning tree (b), in this case, is in fact feasible for the original problem.

5 Conclusions and Future Work

In this paper, we presented a framework for incorporating dynamic cut generation (outer methods) and traditional decomposition methods (inner methods). We have also introduced a new paradigm for the generation of improving inequalities based on decomposition and the separation of solutions to a relaxation, a problem that is often much easier than that of separating arbitrary real vectors. Viewing the cutting plane method, Lagrangian relaxation, and Dantzig-Wolfe decomposition in a common algorithmic framework can yield new insight




Figure 9: Example of a decomposition into minimum spanning trees.

into all three methods. The next step in this research is to complete a computational study that will allow practitioners to make intelligent choices between the the many possible variants we have discussed. As part of this study, we are implementing a generic framework that will allow users to test these methods simply by providing a relaxation, a solver for that relaxation, and separation routines for solutions to the relaxation. Such a framework will enable access to a wide range of alternatives for computing bounds using decomposition.

References

- [1] K. Aardal and S. van Hoesel. Polyhedral techniques in combinatorial optimization. *Statistica Neerlandica*, 50:3–26, 1996.
- [2] R. Ahuja, T. Magnanti, and J. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, Englewood Cliffs, NJ, 1993.
- [3] D. Applegate, R. Bixby, V. Chvátal, and W. Cook. TSP cuts which do not conform to the template paradigm. In *Computational Combinatorial Optimization*, pages 261–303. Springer, 2001.
- [4] P. Augerat, J. M. Belenguer, E. Benavent, A. Corberán, D. Naddef, and G. Rinaldi. Computational Results with a Branch and Cut Code for the Capacitated Vehicle Routing Problem. Technical Report 949-M, Université Joseph Fourier, Grenoble, France, 1995.
- [5] E. Balas and N. Christofides. A restricted Lagrangean approach to the traveling salesman problem. *Mathematical Programming*, 21:19–46, 1981.
- [6] E. Balas and L. Qi. Linear-time separation algorithms for the three-index assignment polytope. *Discrete Applied Mathematics*, 43:1–12, 1993.

- [7] E. Balas and M.J. Saltzman. Facets of the three-index assignment polytope. *Discrete Applied Mathematics*, 23:201–229, 1989.
- [8] E. Balas and M.J. Saltzman. An algorithm for the three-index assignment problem. *Operations Research*, 39:150–161, 1991.
- [9] F. Barahona and R. Anbil. The volume algorithm : Producing primal solutions with a subgradient method. *Mathematical Programming*, 87:385–399, 2000.
- [10] F. Barahona and D. Jensen. Plant location with minimum inventory. *Mathematical Programming*, 83:101–111, 1998.
- [11] C. Barnhart, C. A. Hane, and P. H. Vance. Using branch-and-price-and-cut to solve origin-destination integer multi-commodity flow problems. *Operations Research*, 48:318–326, 2000.
- [12] C. Barnhart, E. L. Johnson, G. L. Nemhauser, M. W. P. Savelsbergh, and P. H. Vance. Branch and price: Column generation for solving huge integer programs. *Operations Research*, 46:316–329, 1998.
- [13] J.E. Beasley. A sst-based algorithm for the steiner problem in graphs. *Networks*, 19:1–16, 1989.
- [14] J.E. Beasley. Lagrangean relaxation. In C.R. Reeves, editor, *Modern Heuristic Techniques for Combinatorial Optimization*. Wiley, 1993.
- [15] A. Belloni and A. Lucena. A Lagrangian heuristic for the linear ordering problem. In M.G.C. Resende and J. Pinho de Sousa, editors, *Metaheuristics: Computer Decision-Making*, pages 123–151. Kluwer Academic, 2003.
- [16] F. C. Calheiros, A. Lucena, and C. de Souza. Optimal rectangular partitions. *Networks*, 41:51–67, 2003.
- [17] A. Caprara, M. Fischetti, and P. Toth. Algorithms for the set covering problem. *Annals of Operations Research*, 98:353–371, 2000.
- [18] P. Carraresi, A. Frangioni, and M. Nonato. Applying bundle methods to optimization of polyhedral functions: An applications-oriented development. Technical Report TR-96-17, Università di Pisa, 1996.
- [19] G. Dantzig and P. Wolfe. Decomposition principle for linear programs. *Operations Research*, 8:101–111, 1960.
- [20] G.B. Danzig and R.H. Ramser. The truck dispatching problem. *Management Science*, 6:80–91, 1959.
- [21] M. Poggi de Aragão and E. Uchoa. Integer program reformulation for robust branch-and-cut-and-price. Working paper, 2004. Available at <http://www.inf.puc-rio.br/~uchoa/doc/rbcp-a.pdf>.

- [22] M.L. Fisher. The Lagrangian relaxation method for solving integer programming problems. *Management Science*, 27:1–18, 1981.
- [23] M.L. Fisher. Optimal solution of vehicle routing problems using minimum k-trees. *Operations Research*, 42:626–642, 1994.
- [24] R. Fukasawa, M. Poggi de Aragão, M. Reis, and E. Uchoa. Robust branch-and-cut-and-price for the capacitated minimum spanning tree problem. In *Proceedings of the International Network Optimization Conference*, pages 231–236, Evry, France, 2003.
- [25] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, New York, 1979.
- [26] A. Geoffrion. Lagrangian relaxation for integer programming. *Mathematical Programming Study*, 2:82–114, 1974.
- [27] M.X. Goemans. The steiner tree polytope and related polyhedra. *Mathematical Programming*, 63:157–182, 1994.
- [28] J. L. Goffin and J. P. Vial. Convex nondifferentiable optimization: A survey focused on the analytic center cutting plane method. Technical Report 99.02, Logilab, Geneva, Switzerland, 1999.
- [29] J. Gondzio and R. Sarkissian. Column generation with a primal-dual method. Technical report, University of Geneva, Logilab, HEC Geneva, 1996.
- [30] Grötschel, M. and Lovász, L. and Schrijver, A. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1:169–197, 1981.
- [31] M. Guignard. Efficient cuts in Lagrangean “relax-and-cut” schemes. *European Journal of Operational Research*, 105:216–223, 1998.
- [32] M. Guignard. Lagrangean relaxation. *Top*, 11:151–228, 2003.
- [33] D. Huisman, R. Jans, M. Peeters, and A. Wagelmans. Combining column generation and Lagrangian relaxation. Technical report, Erasmus Research Institute of Management, Rotterdam, The Netherlands, 2003.
- [34] M. Hunting, U. Faigle, and W. Kern. A Lagrangian relaxation approach to the edge-weighted clique problem. *European Journal of Operational Research*, 131:119–131, 2001.
- [35] A.N. Letchford J. Lysgaard and R.W. Eglese. A new branch-and-cut algorithm for the capacitated vehicle routing problem. *Mathematical Programming*, 100:423–445, 2004.
- [36] N. Kohl, J. Desrosiers, O.G.B. Madsen, M.M. Solomon, and F. Soumis. 2-path cuts for the vehicle routing problem with time windows. *Transportation Science*, 33:101–116, 1999.

- [37] L. Kopman. *A New Generic Separation Routine and Its Application In a Branch and Cut Algorithm for the Capacitated Vehicle Routing Problem*. PhD thesis, Cornell University, May 1999.
- [38] G. Laporte, Y. Nobert, and M. Desrochers. Optimal routing with capacity and distance restrictions. *Operations Research*, 33:1050–1073, 1985.
- [39] J. Linderoth and M. Galati. Knapsack constrained circuit problem. Unpublished working paper, 2004.
- [40] M. E. Lübbecke and J. Desrosiers. Selected topics in column generation. Working paper, Cahier du Gerad, 2002.
- [41] A. Lucena. Tight bounds for the steiner problem in graphs. Talk given at TIMS-XXX - SOBRAPO XXIII Joint International Meeting, Rio de Janeiro, July 1991.
- [42] A. Lucena. Steiner problem in graphs: Lagrangean relaxation and cutting planes. *COAL Bulletin*, 28:2–8, 1992.
- [43] A. Lucena. Non-delayed relax-and-cut algorithms. Working paper, Departamento de Administração, Universidade Federal do Rio de Janeiro, 2004.
- [44] F. Margot, A. Prodon, and T.M. Liebling. Tree polyhedron on 2-tree. *Mathematical Programming*, 63:183–192, 1994.
- [45] C. Martinhon, A. Lucena, and N. Maculan. Stronger k-tree relaxations for the vehicle routing problem. To appear in the European Journal of Operations Research, 2001.
- [46] D. Miller. A matching based exact algorithm for capacitated vehicle routing problems. *ORSA Journal on Computing*, 7:1–9, 1995.
- [47] M. Müller-Hannemann and A. Schwartz. Implementing weighted b-matching algorithms: Towards a flexible software design. In *Proceedings of the Workshop on Algorithm Engineering and Experimentation (ALENEX99)*, volume 1619 of *Lecture notes in computer science*, pages 18–36, Baltimore, MD, 1999. Springer-Verlag.
- [48] D. Naddef and G. Rinaldi. Branch-and-cut algorithms for the capacitated VRP. In P. Toth and D. Vigo, editors, *The Vehicle Routing Problem*, pages 53–84. SIAM, 2002.
- [49] G.L. Nemhauser and L.A. Wolsey. *Integer and Combinatorial Optimization*. Wiley, New York, 1988.
- [50] R.G. Parker and R.L. Rardin. *Discrete Optimization*. Academic Press, Inc., 1988.
- [51] A. Pigatti. *Modelos e Algoritmos para o Problema de Alocação Generalizada e Aplicações*. PhD thesis, Pontifícia Universidade Católica do Rio de Janeiro, 2003.
- [52] R. Prim. Shortest Connection Networks and Some Generalizations. *Bell System Technical Journal*, 36:1389–1401, 1957.

- [53] T.K. Ralphs. *Parallel Branch and Cut for Vehicle Routing*. PhD thesis, Cornell University, May 1995.
- [54] T.K. Ralphs, L. Kopman, W.R. Pulleyblank, and L.E. Trotter Jr. On the capacitated vehicle routing problem. *Mathematical Programming*, 94:343–359, 2003.
- [55] L. M. Rousseau, M. Gendreau, and D. Feillet. Interior point stabilization for column generation. Working paper, Cahier du Gerad, 2002.
- [56] V. Chvátal. *Linear Programming*. W.H. Freeman and Company, 1983.
- [57] J.M. van den Akker, C.A.J. Hurkens, and M.W.P. Savelsbergh. Time-indexed formulations for machine scheduling problems: Column generation. *INFORMS Journal on Computing*, 12:111–124, 2000.
- [58] F. Vanderbeck. Lot-sizing with start-up times. *Management Science*, 44:1409–1425, 1998.
- [59] F. Vanderbeck. A generic view of Dantzig-Wolfe decomposition in mixed integer programming. Working paper, 2003. Available at <http://www.math.u-bordeaux.fr/~fv/papers/dwcPap.pdf>.
- [60] G. Wei and G. Yu. An improved $o(n^2 \log n)$ algorithm for the degree-constrained minimum k-tree problem. Technical report, The University of Texas at Austin, Center for Management of Operations and Logistics, 1995.