

Global optimality conditions for discrete and nonconvex optimization—With applications to Lagrangian heuristics and column generation

Torbjörn Larsson* and Michael Patriksson†

Second revised version submitted January 3, 2005

Abstract

The well-known and established global optimality conditions based on the Lagrangian formulation of an optimization problem are consistent if and only if the duality gap is zero. We develop a set of global optimality conditions that are structurally similar but are consistent for any size of the duality gap. This system characterizes a primal–dual optimal solution by means of primal and dual feasibility, primal Lagrangian ε -optimality, and, in the presence of inequality constraints, a relaxed complementarity condition analogously called δ -complementarity. The total size $\varepsilon + \delta$ of those two perturbations equals the size of the duality gap at an optimal solution. Further, the characterization is equivalent to a near-saddle point condition which generalizes the classic saddle point characterization of a primal–dual optimal solution in convex programming. The system developed can be used to explain, to a large degree, when and why Lagrangian heuristics for discrete optimization are successful in reaching near-optimal solutions. Further, experiments on a set covering problem illustrate how the new optimality conditions can be utilized as a foundation for the construction of new Lagrangian heuristics. Finally, we outline possible uses of the optimality conditions in column generation algorithms and in the construction of core problems.

Subject classification: Programming, integer, theory: global optimality conditions. Programming, integer, algorithms: Lagrangian heuristics, column generation algorithms, core problems.

*Department of mathematics, Linköping University, SE-581 83 Linköping, Sweden. E-mail: tolar@mai.liu.se

†Department of mathematics, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden. E-mail: mi-pat@math.chalmers.se

1 Introduction

1.1 Background

Classic optimality conditions for optimization problems are based on the fulfillment of primal–dual feasibility, primal Lagrangian stationarity or optimality, and complementarity conditions, associated with a particular Lagrangian function. In the convex case, their fulfillment is associated with the compliance of a saddle-point condition for this function.

These conditions are also the foundation for deriving many algorithmic approaches for the search of (near-)optimal solutions. In one such class, the optimality conditions are approximated with simpler systems, such as in sequential quadratic (SQP) or linear (SLP) programming, and interior point methods. Other approaches are associated with satisfying a subset of the conditions while making adjustments in the primal and/or dual spaces in order to satisfy the rest; among these approaches we may count the simplex and dual cutting plane methods for linear programming.

When there is a positive duality gap, however, which is typically the case with integer and combinatorial optimization problems, the primal–dual system describing the set of saddle points of the Lagrangian function is inconsistent. Our first goal is to develop a system of global optimality conditions that is structurally similar, but which is valid for *any* finite duality gap.

1.2 Problem statement and preliminaries

We consider the problem of finding

$$f^* := \text{infimum } f(x), \tag{1a}$$

$$\text{subject to } g(x) \leq 0^m, \tag{1b}$$

$$x \in X, \tag{1c}$$

where the functions $f : \mathbb{R}^n \mapsto \mathbb{R}$ and $g : \mathbb{R}^n \mapsto \mathbb{R}^m$, and the set $X \subseteq \mathbb{R}^n$ is non-empty and closed.

Let

$$\theta(u) := \text{infimum}_{x \in X} \{f(x) + u^T g(x)\}, \quad u \in \mathbb{R}^m, \tag{2}$$

be the dual function associated with the Lagrangian relaxation of the constraints in (1b), and

$$\theta^* := \text{supremum } \theta(u), \tag{3a}$$

$$\text{subject to } u \in \mathbb{R}_+^m, \tag{3b}$$

be the Lagrangian dual problem. The duality gap for this primal–dual pair then is $\Gamma := f^* - \theta^*$.

The applicability of the above-mentioned classic system of optimality conditions does not rely on any additional properties of the problem data. The global optimality and saddle-point conditions are always equivalent; one can, however, not assert that any of them is consistent, without introducing additional properties of problem (1). In order to infer the existence of an optimal solution to the *primal* problem (1) it is enough to assume that f and g are continuous functions and that X is non-empty and bounded; cf. Weierstrass' Theorem. At the same time, both f^* and θ^* are finite, and therefore also the size of the duality gap Γ . In order to infer the existence of an optimal solution to the Lagrangian *dual* problem (3) we must add regularity properties of problem (1), and to ensure that the duality gap is zero (so that the global optimality conditions mentioned earlier are consistent) we need to further impose a convexity property. For further relations between the primal–dual problem and the classic system of optimality conditions, see, e.g., [BSS93, Ber99, BNO03].

Also our global optimality conditions, to be developed in the next section, are applicable without any additional properties of the problem data. (In any case, the reader may assume throughout our presentation that there exist feasible solutions to problem (1) and that enough properties are imposed such that problem (3) also is well-defined.) The optimal solution set of problem (1) is denoted X^* , a member of which is denoted x^* ; an ε -*optimal* solution to any problem refers to any vector that is feasible and deviates in objective value at most ε from the optimal one; the set of such vectors in problem (1) is denoted X^ε .

Given $(x, u) \in X \times \mathbb{R}_+^m$, we define the classic *global optimality conditions* for the problem (1) referred

to above as the combination of Lagrangian optimality, primal feasibility, and complementarity:

$$f(x) + u^T g(x) \leq \theta(u), \tag{4a}$$

$$g(x) \leq 0^m, \tag{4b}$$

$$u^T g(x) = 0. \tag{4c}$$

The following result establishes the consistency of the system (4). Similar results can be found in [Sha79a, Theorem 5.1] and [BSS93, Theorem 6.2.5].

THEOREM 1 (primal–dual optimality conditions). *Let $(x, u) \in X \times \mathbb{R}_+^m$. Then, the following three statements are equivalent.*

- (i) *The pair (x, u) satisfies the system (4).*
- (ii) *The pair (x, u) is a saddle point of the Lagrangian function $(x, u) \mapsto L(x, u) := f(x) + u^T g(x)$ over $X \times \mathbb{R}_+^m$, that is,*

$$L(x, v) \leq L(x, u) \leq L(y, u), \quad \forall (y, v) \in X \times \mathbb{R}_+^m.$$

- (iii) *x solves the primal problem (1), u solves the dual problem (3), and $f^* = \theta^*$. □*

COROLLARY 2 (a characterization of optimality). *Given any $u \in \mathbb{R}_+^m$,*

$$\{x \in X \mid (4) \text{ is satisfied}\} = \begin{cases} X^*, & \text{if } \theta(u) = f^*, \\ \emptyset, & \text{if } \theta(u) < f^* \end{cases}$$

holds. □

There are two possible cases in which $\theta(u) < f^*$, whence the system (4) is inconsistent: (a) the vector u is not optimal in the Lagrangian dual problem (3); (b) there is a positive duality gap, that is, $\theta^* < f^*$. Convex problems satisfying a constraint qualification have a zero duality gap (cf. [BSS93, Theorem 6.2.4], [Ber99, Chapter 5], and [BNO03, Chapter 6]); linear programming is another important class of such problems. In the nonconvex case, Corollary 2 in general becomes a theorem about the empty set. Note further the following easy consequences of the above, which are quite negative results: if there is a positive duality gap, then a primal optimal solution will *not* satisfy the global optimality conditions (4) for *any* value of $u \geq 0^m$; further, if the relaxed constraints are equalities (see Section 2.2 for an analysis), then the solution to the Lagrangian relaxed problem can, in the case of a positive duality gap, never yield a primal optimal solution, regardless of the choice of dual vector.

Nevertheless, *Lagrangian relaxation* is a popular and often successful approach to integer and combinatorial optimization problems, which is somewhat surprising given that it is in fact also based on the system (4). The interest in Lagrangian relaxation started with the publication of the papers by Held and Karp [HeK70, HeK71] on the application of subgradient optimization to the Lagrangian dual of a formulation of the traveling salesman problem, and grew with the subsequent publications in the 1970's and early 1980's (e.g., [Las70, HWC74, Geo74, MSW76, Erl78, Sha79a, Sha79b, BGB81, Fis85]); see also the survey papers [Fis81, Bea93]. Until the 1980's, Lagrangian relaxation was often used as a bounding procedure in branch and bound algorithms (cf. [Fis81]). Since then, interest in Lagrangian relaxation and solution techniques for (approximately) maximizing the dual function, has shifted towards their combination with a *Lagrangian heuristic* (cf. [Fis81] and [Bea93, Section 6.4]), that adjusts the (typically) infeasible Lagrangian optimal solutions into feasible ones.

To be more precise, we shall define a Lagrangian heuristic as follows: *Initiated at a primal vector in the set defined by the non-relaxed constraints, it adjusts this vector by (a) executing a finite number of operations and is further characterized by (b) its utilization of information from the Lagrangian dual problem, (c) the sequence of vectors generated remaining within the set of non-relaxed constraints, and (d) the terminal vector being, if possible, primal feasible and hopefully also near-optimal in the problem (1).*

Included in this definition is the possibility that the heuristic does not terminate at the first primal feasible solution found, but continues with a primal local search. We remark that the most common means in which to comply with the property (b) is to initiate the heuristic at a (near-)optimal solution

to a Lagrangian relaxed problem or to perform the adjustments guided by a merit function defined by the Lagrangian cost.

We distinguish between two types of Lagrangian heuristics that we will analyze. The first, which will be referred to as *conservative*, has the properties that the initial vector is a (near-)optimal Lagrangian problem solution, that the moves are local only in the sense that each move changes the Lagrangian objective function value slightly, and that the number of local moves that may be performed normally is required to be small; consequently, a conservative Lagrangian heuristic will retain near-optimality in the Lagrangian problem. A conservative heuristic may provide anything from very good primal feasible solutions to no feasible solutions at all, depending on the nature of the problem being attacked by Lagrangian relaxation and the design of the heuristic.

The second type of Lagrangian heuristic is non-conservative, and will therefore be referred to as *radical*. It has the property that it allows the resulting primal vector to be far from optimal in the Lagrangian relaxed problem. This category includes heuristics that are defined by the solution of a restriction of the original problem (such as the Benders subproblem), and large scale neighbourhood search; a more detailed description of them are given in Section 3.2. A conservative heuristic will typically turn into a radical one if it is allowed to run forever.

In the next section we develop a set of novel optimality conditions for the general problem (1). These conditions are structurally similar to those in the convex case, but are valid for any size of the duality gap, including zero [in which case the new conditions reduce to (4)]. The new conditions provide an understanding of the success or failure of conservative Lagrangian heuristics: a conservative heuristic will fail if the duality gap is large (because no feasible solution is found), but it will also fail when the duality gap is small if the design of the heuristic ignores that near-optimal solutions are near-complementary (in which case a feasible solution may not be near-optimal). Our analysis also shows that if the duality gap is large, then in order to be able to find near-optimal solutions, it is necessary to consider radical heuristics that allow the primal solution to significantly deviate both from Lagrangian optimality and complementarity fulfillment. The importance of approximate complementarity fulfillment has to our knowledge never before been considered in Lagrangian heuristics.

2 Global optimality conditions

The development of a new set of global optimality conditions for problem (1) is a primary goal of this paper. In contrast to (4), the new system contains perfect information about the primal–dual set of optimal solutions regardless of the size of the duality gap, and it lends itself very well to the analysis and construction of Lagrangian heuristics. Since system (4) is not consistent for problems with a positive duality gap, it seems reasonable to investigate *relaxations* of it.

2.1 Inequality constraints

We introduce nonnegative numbers ε and δ . Given the pair $(x, u) \in X \times \mathbb{R}_+^m$, we define the *global optimality conditions* for the problem (1) as

$$f(x) + u^T g(x) \leq \theta(u) + \varepsilon, \tag{5a}$$

$$g(x) \leq 0^m, \tag{5b}$$

$$u^T g(x) \geq -\delta, \tag{5c}$$

$$\varepsilon + \delta \leq \Gamma, \tag{5d}$$

$$\varepsilon, \delta \geq 0. \tag{5e}$$

In this system, (5a) and (5c) define ε -optimality in the Lagrangian problem, and δ -complementarity, respectively. The systems (4) and (5) are equivalent precisely when the duality gap is zero. Note that in the system (5) the size Γ of the duality gap is present explicitly.

The following theorem provides the analogous result to Theorem 1.

THEOREM 3 (primal–dual optimality conditions). *Let $(x, u) \in X \times \mathbb{R}_+^m$. Then, the following three statements are equivalent.*

- (i) *Together with the pair (ε, δ) , the pair (x, u) satisfies the system (5).*

(ii) The pair $(\varepsilon, \delta) \geq (0, 0)$ satisfies $\varepsilon + \delta = \Gamma$, and the pair (x, u) satisfies the following saddle-point like condition for the Lagrangian function $(x, u) \mapsto L(x, u)$ over $X \times \mathbb{R}_+^m$:

$$L(x, v) - \delta \leq L(x, u) \leq L(y, u) + \varepsilon, \quad \forall (y, v) \in X \times \mathbb{R}_+^m. \quad (6)$$

(iii) x solves the primal problem (1) and u solves the dual problem (3).

PROOF. We establish first that (i) and (ii) are equivalent. It is clear that the consistency of (5) implies that $\varepsilon + \delta = \Gamma$ holds [combine (5a), (5c), and the duality gap consequence that $f(x) - \theta(u) \geq \Gamma$ holds with (5d)]. That the second inequality in (6) is equivalent to (5a) is immediate. The first inequality in (6) is equivalent to

$$g(x)^\top(u - v) \geq -\delta, \quad \forall v \in \mathbb{R}_+^m. \quad (7)$$

With $v = 0^m$, we obtain (5c). To reach (5b), we note that if it is not satisfied, then there is some $i \in \{1, 2, \dots, m\}$ for which $g_i(x) > 0$; by letting $v_i \rightarrow +\infty$, we contradict (7). Conversely, for all $v \in \mathbb{R}_+^m$,

$$\begin{aligned} f(x) + v^\top g(x) - \delta &= f(x) + u^\top g(x) + g(x)^\top(v - u) - \delta \\ &\leq f(x) + u^\top g(x), \end{aligned}$$

where the inequality follows from (5b)–(5c). This completes the first part of the proof.

Next, we establish that (i) and (iii) are equivalent. Suppose that (i) holds. Then, (5a), (5c), and (5d) imply that

$$f(x) \leq \theta(u) + \Gamma. \quad (8)$$

By definition, $\Gamma = f^* - \theta^*$. Therefore, (5) holds if and only if (x, u) is primal–dual optimal, whence (iii) follows.

Finally, suppose that (iii) holds. Then, (8) holds. Further, suppose that for the given pair (x, u) , we choose ε and δ according to

$$\varepsilon := \varepsilon(x, u) = \Gamma + u^\top g(x) \quad \text{and} \quad \delta := \delta(x, u) = -u^\top g(x). \quad (9)$$

Adding $u^\top g(x)$ to both sides of the inequality (8) yields (5a). The inequality (5b) follows from the optimality of x in the primal problem (1), and (5c) is trivially satisfied, by the choice (9), and (i) follows. This completes the proof. \square

Theorem 3 implies the following (cf. Corollary 2):

COROLLARY 4 (a characterization of optimality). *Given any $u \in \mathbb{R}_+^m$,*

$$\{x \in X \mid (5) \text{ is satisfied}\} = \begin{cases} X^*, & \text{if } \theta(u) = f^* - \Gamma, \\ \emptyset, & \text{if } \theta(u) < f^* - \Gamma \end{cases}$$

holds. \square

The systems (4) and (5) are structurally similar characterizations of the (Cartesian product) sets of primal–dual optimal solutions. The main difference between the convex and nonconvex cases is that, in the latter, primal–dual optimal solutions are *near-saddle points*; no saddle points exist when $\Gamma > 0$. An interpretation is given in Figure 1 in the context of the example of Section 2.3.

In the case of convex programming, the global optimality conditions (5) are related to, but much simpler than, those presented by Strodiot et al. [SNH83], which are based on an ε -subdifferential form of a global Karush–Kuhn–Tucker condition.

If the Lagrangian relaxed problem (2) separates into a number of independent problems, then the value of the Lagrangian optimality perturbation ε in condition (5a) is the sum of perturbations in the independent Lagrangian problems; cf. the definition of ε in (24). In such cases, condition (5a) can be similarly disaggregated into a set of conditions, one for each independent problem. Condition (5c) can *always* be disaggregated into multiple conditions, at most one for each relaxed constraint; the value of perturbation δ is then the sum of individual perturbations.

We next present a relaxation of the system (5), which is consistent also for near-optimal solutions in the primal and dual problems. We will use this system particularly when analyzing algorithms, such as

Lagrangian heuristics, for integer programs later on in this paper. To this end, we introduce a nonnegative parameter κ , which defines the level of primal–dual near-optimality allowed.

Given the pair $(x, u) \in X \times \mathbb{R}_+^m$, we define the *relaxed global optimality conditions* for the problem (1) as

$$f(x) + u^\top g(x) \leq \theta(u) + \varepsilon, \quad (10a)$$

$$g(x) \leq 0^m, \quad (10b)$$

$$u^\top g(x) \geq -\delta, \quad (10c)$$

$$\varepsilon + \delta \leq \Gamma + \kappa, \quad (10d)$$

$$\varepsilon, \delta, \kappa \geq 0. \quad (10e)$$

We note immediately, with reference to the above result, that a consistent system (10) always has

$$\Gamma \leq \varepsilon + \delta \leq \Gamma + \kappa.$$

For this system, we state some immediate consequences in terms of the relations to near-optimal, and near-complementary solutions to the primal–dual problem.

PROPOSITION 5 (near-optimal solutions and the system (10)).

- (a) (near-optimality in the primal problem (1)). *Let $(x, u) \in X \times \mathbb{R}_+^m$. Suppose that, for some $\varepsilon, \delta, \kappa \geq 0$, (10) holds. Then, x is feasible in (1), and*

$$f(x) \leq \theta(u) + \Gamma + \kappa.$$

Suppose further that u solves the dual problem (3). Then,

$$f(x) \leq f^* + \kappa.$$

- (b) (near-optimality in the Lagrangian problem (2)). *Suppose that $(x, u) \in X \times \mathbb{R}_+^m$ is β -optimal and α -optimal, respectively, in the primal and dual problem (1) and (3), for some $\beta, \alpha \geq 0$. Then,*

$$\theta(u) \leq f(x) + u^\top g(x) \leq \theta(u) + \Gamma + \beta + \alpha.$$

Suppose further that $(x, u) \in X \times \mathbb{R}_+^m$ solves the primal and dual problems (1) and (3), respectively. Then,

$$\theta^* \leq f(x) + u^\top g(x) \leq f^*. \quad (11)$$

- (c) (near-complementarity). *Suppose that $(x, u) \in X \times \mathbb{R}_+^m$ is β -optimal and α -optimal, respectively, in the primal and dual problems (1) and (3), for some $\beta, \alpha \geq 0$. Suppose further that $\varepsilon \geq 0$ is such that (10a) holds with equality. Then, (10) holds, with $\delta := \Gamma - \varepsilon + \alpha + \beta \geq 0$ and $\kappa := \alpha + \beta$. In fact,*

$$\Gamma + \beta + \alpha \leq \varepsilon + \delta \leq \Gamma + \kappa$$

always holds when (10) is consistent.

Suppose further that $(x, u) \in X \times \mathbb{R}_+^m$ solves the primal and dual problems (1) and (3), respectively. Then, (10) holds, with $\delta := \Gamma - \varepsilon \geq 0$ and $\kappa = 0$. \square

REMARK 6 (interpretations). The result in (a) states that vectors x that are near-optimal in Lagrangian problem (2) and near-complementary also are near-optimal solutions to primal problem (1), in particular, when the value of κ is small (that is, when the sum of the perturbations ε and δ are in the order of the size of the duality gap). Result (a) implies that the goal, when searching for a primal vector x that satisfies (5), should be to minimize κ , that is, essentially minimizing $\varepsilon + \delta$. A specialization of result (a) to linear integer programming is found in [NeW88, Corollary II.3.6.9].

The result in (b) shows that a (near-)optimal solution to primal problem (1) must also be near-optimal in the Lagrangian problem defined at a (near-)optimal dual solution. The example in Section 2.3 will show that either of (or neither of) the two inequalities in (11) may be tight for some optimal solutions. (In the case of equality constraints, however, the last inequality is always tight at optimal solutions.)

The result in (c) shows that a (near-)optimal solution to the primal problem (1) must also be near-complementary. It shows how closely related the two perturbations ε and δ are to the value of Γ , and it follows that system (5) is always consistent at an optimal primal–dual solution. \square

Theorem 3 also implies the following:

COROLLARY 7 (a characterization of near-optimality). *Let $u \in \mathbb{R}_+^m$ be α -optimal in the dual problem (3), for some $\alpha \geq 0$. Then,*

$$\{x \in X \mid (10) \text{ is satisfied}\} = \begin{cases} X^{\kappa-\alpha}, & \text{if } \kappa \geq \alpha, \\ \emptyset, & \text{if } \kappa < \alpha \end{cases} \quad (12)$$

holds. □

REMARK 8 (interpretations). We characterize the optimal solution set X^* precisely when $\kappa = \alpha$. From this characterization, we see that primal optimal solutions can be obtained from non-optimal dual solutions, provided that the sum of the perturbations ε and δ matches precisely this non-optimality, in the sense that $\varepsilon + \delta = \Gamma + \kappa = \Gamma + \alpha$. (The sum is clearly unique, but not necessarily the values of ε and δ individually; cf. the example in Section 2.3.) Since this result is more general than Corollary 2, it follows that it is true also for convex problems.

Theorem 1 and Corollary 2 are special cases of Theorem 3 and Corollary 7, and follow when in addition $\Gamma = \varepsilon = \delta = \kappa = \alpha = 0$. Moreover, Corollary 4 is the special case of the above, for the case where $\kappa = \alpha = 0$.

The relation (12) combines the results of the above theorem with that of Proposition 5, in that we relate near-optimal solutions of the primal problem (1) and the dual problem (2) to each other. [One can of course also characterize the set of *dual* (near-)optimal solutions through a statement analogous to (12).] □

A special case of Corollary 7 related to linear integer programming is found in [NeW88, Theorem II.3.6.7], however stated in terms of objective values. (That theorem traces back to results in Everett [Eve63], and is outlined in [Las70, Section 8.3.2].)

REMARK 9 (on the use of the system (5)). If the set X is discrete, then it is possible, in principle, to solve problem (1) by enumerating the points in X according to an increasing value of the Lagrangian function $(x, u) \mapsto f(x) + u^T g(x)$. Every time a feasible solution appears, we obtain an upper bound on f^* . It is easy to show that, during the enumeration, the value of the Lagrangian function always underestimates the objective value of every feasible solution that still has *not* been found. In particular, if the enumeration continues until the value of the Lagrangian function becomes at least as large as the best known upper bound corresponding to a feasible solution, then that corresponding primal vector x is globally optimal. (As a special case, with the choice of $u = 0^m$, we recover the simple method in which the first feasible solution found is optimal, as then the enumeration is made in terms of the original cost.) If, on the other hand, the enumeration is terminated prior to this occurrence, then the terminal value of the Lagrangian function is a lower bound to f^* .

Handler and Zang [HaZ80] utilized a ranking methodology based on the Lagrangian cost to solve a knapsack-constrained shortest-path problem, from an optimal dual solution u . A recent contribution to this methodology is found in Carlyle and Wood [CaW03]. Recently, Caprara et al. [CFT02] have constructed a similar feasibility heuristic for a train timetabling problem. A similar methodology with the purpose of constructing core problems is presented in Section 5.1, for the case where X is a discrete Cartesian product set. □

2.2 Equality constraints

We next specialize the above main result to the case of equality constraints. So, suppose, locally in this section only, that (1b) is replaced by

$$h(x) = 0^\ell, \quad (1b)$$

where $h : \mathbb{R}^n \mapsto \mathbb{R}^\ell$ is continuous. The multiplier vector for these constraints is $v \in \mathbb{R}^\ell$; the dual function $\theta : \mathbb{R}^\ell \mapsto \mathbb{R}$ is defined accordingly. The condition corresponding to (5) then is

$$f(x) + v^T h(x) \leq \theta(v) + \varepsilon, \quad (13a)$$

$$h(x) = 0^\ell, \quad (13b)$$

$$0 \leq \varepsilon \leq \Gamma. \quad (13c)$$

THEOREM 10 (primal–dual optimality conditions). *Let $(x, v) \in X \times \mathbb{R}^\ell$. Then, the following two statements are equivalent.*

- (i) *Together with ε , the pair (x, v) satisfies the system (13).*
- (ii) *The perturbation $\varepsilon = \Gamma$, and the pair (x, v) satisfies the following saddle-point like condition for the Lagrangian function $(x, v) \mapsto L(x, v) := f(x) + v^\top h(x)$ over $X \times \mathbb{R}^\ell$:*

$$L(x, w) \leq L(x, v) \leq L(y, v) + \varepsilon, \quad \forall (y, w) \in X \times \mathbb{R}^\ell.$$

- (iii) *x solves the primal problem (1) and v solves the dual problem (3).* □

COROLLARY 11 (a characterization of optimality). *Given any $v \in \mathbb{R}^\ell$,*

$$\{x \in X \mid (13) \text{ is satisfied}\} = \begin{cases} X^*, & \text{if } \theta(v) = f^* - \Gamma, \\ \emptyset, & \text{if } \theta(v) < f^* - \Gamma \end{cases}$$

holds. □

The relaxed optimality conditions here are:

$$f(x) + v^\top h(x) \leq \theta(v) + \varepsilon, \tag{14a}$$

$$h(x) = 0^\ell, \tag{14b}$$

$$\varepsilon \leq \Gamma + \kappa, \tag{14c}$$

$$\varepsilon, \kappa \geq 0. \tag{14d}$$

COROLLARY 12 (a characterization of near-optimality). *Let $v \in \mathbb{R}^\ell$ be α -optimal in the dual problem (3), for some $\alpha \geq 0$. Then,*

$$\{x \in X \mid (14) \text{ is satisfied}\} = \begin{cases} X^{\kappa-\alpha}, & \text{if } \kappa \geq \alpha, \\ \emptyset & \text{if } \kappa < \alpha \end{cases}$$

holds. □

2.3 A numerical example

Consider the following linear integer programming problem:

$$f^* := \text{minimum } f(x) := -x_2, \tag{15a}$$

$$\text{subject to } g(x) := x_1 + 4x_2 - 6 \leq 0, \tag{15b}$$

$$x \in X := \{x \in \mathbb{Z}^2 \mid 0 \leq x_1 \leq 4; 0 \leq x_2 \leq 2\}. \tag{15c}$$

The Lagrangian function associated with the dualization of the constraint (15b) with a multiplier $u \geq 0$ is $(x, u) \mapsto L(x, u) := ux_1 + (4u - 1)x_2 - 6u$. The dual problem has the objective function

$$\theta(u) := \begin{cases} 2u - 2, & 0 \leq u \leq 1/4, \\ -6u, & 1/4 \leq u, \end{cases}$$

whose maximum over \mathbb{R}_+ is attained at $u = 1/4$, with $\theta^* = \theta(u) = -3/2$.

The linear programming (LP) relaxation $[\mathbb{R}^2$ replaces \mathbb{Z}^2 in (15c)] of the above problem has the same Lagrangian dual problem, and its primal solution is characterized by the system (4) as follows. At $u = 1/4$, the minimum of the Lagrangian function over the set $[0, 4] \times [0, 2]$ is the set $X(u) := \{x \in \mathbb{R}^2 \mid x_1 = 0; x_2 \in [0, 2]\}$. [This is the set defined by (4a).] Together with primal feasibility [that is, (4b), or, in this case, (15b)], we obtain that x_2 is further restricted to be less than or equal to $3/2$, while complementarity [that is, (4c)], forces x_2 to take on the value $3/2$. So, from (4) we obtain that the primal–dual optimal solution set is the singleton set $\{(0, 3/2)^\top\} \times \{1/4\}$, with optimal (and saddle) value $-3/2$.

Returning to the integer program, there are three optimal solutions, $x^1 = (0, 1)^T$, $x^2 = (1, 1)^T$, and $x^3 = (2, 1)^T$, with objective value $f^* = -1$. The duality gap is $\Gamma := f^* - \theta^* = 1/2$. In order to show how these optimal solutions will arise from an application of the system (5), we begin by noting that ε and δ must sum to $\Gamma = 1/2$, according to Theorem 3.

We first investigate the case where $\varepsilon = 0$, that is, the Lagrangian problem is solved exactly. Then, from (5a), we obtain that $X(u) = \{(0, 0)^T, (0, 1)^T, (0, 2)^T\}$. Primal feasibility [that is, (5b)] then dictates that x is either $(0, 0)^T$ or $(0, 1)^T$. Finally, we know that $u^T g(x) \geq -1/2$ [that is, (5c)] since $\delta = \Gamma - \varepsilon = 1/2$. The only primal vector of the two satisfying (5) with $\varepsilon = 0$ is $x^1 = (0, 1)^T$. (So, this solution violates the complementarity conditions.)

That the system (5) is consistent when $\varepsilon = 0$ is not necessarily the case; ε may need to take on positive values in order to reach an optimal solution. The optimal solution $x^2 = (1, 1)^T$ corresponds to letting $\varepsilon = \delta = 1/4$. In Figure 1, this optimal solution is contrasted with a non-optimal solution, and the values of ε and δ are given a further interpretation.




Figure 1: Illustration of the role of ε and δ in the characterization of optimality. For the optimal solution x^2 , the value of $\varepsilon(x^2, u^*)$ equals the vertical distance between the two functions θ and $L(x^2, \cdot)$ at u^* . The remaining vertical distance to f^* equals minus the slope of the function $L(x^2, \cdot)$ at u^* [which is $g(x^2) = -1$] times u^* , that is, $\delta(x^2, u^*) = 1/4$. In the case of the candidate vector $\bar{x} := (2, 0)^T$, the value of ε is $1/2$, and $\delta = 1$ [the slope of $L(\bar{x}, \cdot)$ at u^* is -4]; in this case, then, $\theta^* + \varepsilon + \delta = f(\bar{x}) = 0 > f^*$, so \bar{x} cannot be optimal.

The optimal solution $x^3 = (2, 1)^T$ corresponds to letting $\varepsilon = 1/2$, while $\delta = 0$. This solution violates Lagrangian optimality even more than in the previous solution, but it is complementary.

We illustrate the three optimal solutions also in Figure 2, in the same order. In the figure, we have plotted the feasible region of the problem (15) intersected with the vectors x that satisfy Lagrangian ε -optimality and δ -complementarity at u^* for the three values of (ε, δ) given above.

3 A dissection of Lagrangian heuristics

While Lagrangian heuristics are designed mainly to identify primal *feasible* solutions, an important goal is to also reach *near-optimal* solutions. We may interpret a Lagrangian heuristic as a procedure for attempting to satisfy the system (10); according to Theorem 3 and its corollary, the heuristic should, however, also be designed to recover a primal feasible solution x such that the corresponding value of




Figure 2: The three optimal solutions x^1 , x^2 , and x^3 (marked with larger dots) to the problem (15) are specified by the global optimality conditions (5), respectively for $(\varepsilon, \delta) := (0, 1/2)$, $(1/4, 1/4)$, and $(1/2, 0)$. The shaded regions and arrows illustrate the conditions (5a) and (5c) corresponding to $u = u^*$.

κ , and hence of ε and δ , is small, relative to the size of the duality gap which their sum must not underestimate. If, and only if, this is possible, then near-optimal solutions to (1) are identified.

We stress that for *equality* constrained problems, the value of ε is *fixed* at Γ for every optimal solution [statement (ii) in Theorem 10]. Lagrangian heuristics that attempt to keep the value of ε as small as possible while striving for feasibility in the relaxed constraints will therefore typically result in near-optimal solutions for such problems, since making local moves will guarantee that Lagrangian problem near-optimality is retained.

In the case of *inequality* constrained problems, however, only the sum $\varepsilon + \delta$ is fixed, at Γ [statement (ii) in Theorem 3]; the respective sizes of ε and δ may vary significantly, not only among problem instances, but even among optimal solutions for the same problem. (See the example in Section 2.3.) Therefore, whether a Lagrangian heuristic will be successful or not depends on several additional factors whose natures are not so easy to determine in advance. For example, minimizing the value of ε in this context may result in an inconsistent system, thus making the heuristic fail to produce a feasible solution.

This section collects some basic consequences of Theorem 3 in terms of the workings of a successful Lagrangian heuristic, depending on the type of problem being attacked through Lagrangian relaxation.

We begin by describing the connection between the system (10) and Lagrangian heuristics, according to the definition in Section 1. We have at hand some dual vector $u \in \mathbb{R}_+^m$, which is α -optimal in the problem (3) for some (unknown) $\alpha \geq 0$. We attack the Lagrangian problem (2), obtaining a primal solution $\bar{x}(u) \in X$ which is ε_0 -optimal in (2) for some (possibly unknown) $\varepsilon_0 \geq 0$. [Thus, we satisfy (10a), with $\varepsilon = \varepsilon_0$.] For future reference, we also introduce $\delta_0(u) := -u^T g(\bar{x}(u)) \in (-\infty, \infty)$ to denote the level of complementarity fulfillment at $\bar{x}(u)$. If $\bar{x}(u)$ does not satisfy the relaxed constraints, then an attempt is made to attain primal feasibility through a manipulation of this primal solution, while remaining within the set X . Sometimes, this manipulation is not terminated when a primal feasible solution has been found, but is instead followed by a primal local search heuristic. If successful, the result of this heuristic projection of the infeasible solution $\bar{x}(u)$ onto the feasible set of the problem (1) is a vector, \bar{x} . This vector is associated with the values $\varepsilon \geq 0$ and $\delta := -u^T g(\bar{x}) \geq 0$, satisfying (10a) and (10c). The vector \bar{x} is, further, β -optimal in (3), where $\beta \geq 0$ satisfies the relation in Corollary 7. [Obviously, $\beta \leq f(\bar{x}) - \theta(u)$.]

Striving for feasibility while simultaneously fulfilling near-complementarity and Lagrangian near-optimality clearly is strongly related to striving for a feasible solution that is near-optimal (because any primal–dual pair (x, u) which satisfies (10) necessarily fulfills $f(x) \leq \theta(u) + \varepsilon + \delta$). The reason for characterizing Lagrangian heuristics in terms of the fulfillment of near-complementarity and Lagrangian near-optimality, instead of simply considering the primal objective value (which is the immediate choice), is twofold: the appearance of the global optimality conditions makes this viewpoint natural, and it provides an insight into the nature of Lagrangian heuristics. Further, the findings described in the remainder of the paper show that the viewpoint of near-complementarity and Lagrangian near-optimality fulfillment can be used as a means for deriving new solution methods.

3.1 Conservative heuristics

Suppose that Γ , α , and δ_0 all are close to zero. Since $\varepsilon + \delta \leq \Gamma + \kappa$ holds at any solution to the system (10) according to Proposition 5(c), both ε and δ are known to be close to zero at near-optimal primal solutions. The near-optimal primal solutions therefore all lie in a small neighbourhood [in the sense of the value of $L(\cdot, u)$] of the Lagrangian optimal solution $x(u)$; consequently, when $\varepsilon_0 = 0$ holds, it is sufficient to consider a feasibility heuristic which is conservative in the sense of the adjustment in the value of $L(\cdot, u)$.

Conclusion When the dual solution at hand is near-optimal and the duality gap is small, and when the values of ε_0 and δ_0 are small, then in order to be able to find near-optimal primal feasible solutions, it is always *sufficient* to consider Lagrangian heuristics that are conservative in the value of the Lagrangian function.

If the value of ε_0 or δ_0 is large, then it is necessary that feasibility and Lagrangian optimality (and possibly also complementarity fulfillment) are improved simultaneously in the heuristic. (This case might arise, for example, if the Lagrangian relaxed problem is a difficult discrete optimization problem.) The method would therefore then not qualify as a conservative heuristic.

Some examples of conservative heuristics which have yielded very good results for certain combinatorial problems are given in [JLV90, CaM91, Fis94].

A special case of the above is *convex programming*, where $\Gamma = \varepsilon = \delta = 0$ (if $\alpha = 0$). The specific design of a proper Lagrangian heuristic for this case depends largely of the characteristics of the problem in question.

A feasibility heuristic which is not based on a projection-like operation may clearly result in large values of ε , and therefore in solutions of low quality, even when applied at near-optimal dual solutions. To provide an example of this failure, we consider dual solution procedures for strictly convex minimum cost network flow problems. Few articles are devoted to the generation of primal feasible flows in this application; one of them is [Ven91]. From a dual vector u^t , his heuristic works as follows. Unless u^t is optimal, the resulting Lagrangian problem solution, $x(u^t)$, does not satisfy all the flow conservation constraints. A *linear* network flow problem is then constructed, where the demand vector is the residual $b - Ax(u^t)$, $A \in \{-1, 0, 1\}^{m \times n}$ and $b \in \mathbb{R}^m$ being the node-link incidence matrix and the demand vector, respectively, and where the linear cost vector is $\nabla f(x(u^t))$. Letting ξ be a solution to this problem, the vector $x(u^t) + \xi$ is feasible in the original problem. The quality of this solution may however be poor, since the use of a linear cost makes the solution extremal, whereas the original problem will typically have an optimal solution in the relative interior of the feasible set. In the notation of this paper, we conclude that the heuristic is not conservative (ε may be large), since the size of ξ can be substantial even when we are close to a dual optimum (and, further, it does not tend to zero as $\{u^t\} \rightarrow u^*$). This observation lead Marklund [Mar93], in a master's project supervised by the authors, to devise heuristic projections based on conservative node imbalance-reducing graph search techniques, which, in comparison, yield feasible flows of a much better quality. (See further [Pat94, Chapter 4.3] for the only published account of the procedures in [Mar93].)

3.2 Radical heuristics

Suppose that Γ is large, and that α or δ_0 are large also. In the first case, we do not know beforehand whether the identity $\varepsilon + \delta = \Gamma + \alpha$ (assuming that $\beta = 0$) requires ε or δ , or both, to be large, we cannot guarantee that a conservative heuristic (which keeps the value of ε low) will be able to produce near-optimal solutions to (1). (In some cases it *may* be successful, since violating complementarity could compensate for a small value of ε . In the equality constrained case, however, conservative heuristics cannot yield feasible solutions if Γ is large.) This implies that it is necessary to choose a heuristic which is radical in that it allows ε to be large.

If the value of δ is unknown and the value of δ_0 is large, then the heuristic must be radical in order to allow δ to take on either small or large values.

Conclusion When the dual solution at hand is far from being optimal, the duality gap is large, or the initial complementarity violation is large, then, in order to be able to obtain near-optimal primal feasible solutions, it is *necessary* to consider Lagrangian heuristics that are radical with respect to the value of

the Lagrangian function and allow for both small and large complementarity violations. (In the case of equality constraints, the remark on complementarity may be stricken.)

Because of the radical, and therefore global, nature of the above type of heuristics, it may be more appropriate to think of them as being global heuristic optimization procedures, as opposed to the local nature of conservative heuristics. One example of a radical type of heuristic is *very large-scale neighbourhood search* (see [AEOP02]). Another example of a heuristic that may be designed to be radical in our sense is the class of *greedy algorithms* for discrete optimization. Still another technique which may serve as a radical Lagrangian heuristic is available in integer optimization problems with two groups of variables, such as mixed-integer linear programs: the restriction to the original problem known as the *Benders subproblem* (see, e.g., [Las70]). Suppose that $x = (x_1, x_2)$, where the subvector x_1 is required to be integer valued whereas the subvector x_2 may take on fractional values. Such models are frequent in design problems in which the integer variables are associated with design decisions, while the continuous variables are associated, for example, with network flows. The Lagrangian heuristic then is to solve the (linear) Benders subproblem for the original problem, which means that the original objective function is optimized with respect to continuous variables x_2 while integer variables x_1 are fixed to their values $\bar{x}_1(u)$ from the Lagrangian problem. In some cases, the Lagrangian relaxation is such that the resulting solution \bar{x}_1 will be feasible; otherwise, the integer solution may have to be updated heuristically and the Benders subproblem resolved. This heuristic clearly qualifies as a radical heuristic, as it allows for large moves in x_2 , and therefore also in the value of ε . (Depending on the problem instance, however, even a large move in some variables could amount to a small adjustment in the value of ε .) A successful application of this type of heuristic is given in [MuC79] for a location problem.

3.3 A summary

From the above discussions we can now characterize the applicability of Lagrangian heuristics. When the dual solution is near-optimal and the duality gap is small, and ε_0 and δ_0 take on small values, then in order to reach a near-optimal primal feasible solution it is enough to consider the class of conservative heuristics; in all other cases, it is necessary to consider radical Lagrangian heuristics in order to guarantee finding near-optimal solutions.

4 Experiments on the set covering problem

Our goal with the experiments of this section is to illustrate the potential of utilizing our theoretical findings, in particular by considering also complementarity fulfillment in Lagrangian heuristics for such problems.

4.1 The set covering problem and its dual

The set covering problem is to find

$$f^* := \text{minimum} \sum_{j=1}^n c_j x_j, \tag{16a}$$

$$\text{subject to} \sum_{j=1}^n a_j x_j \geq 1^m, \tag{16b}$$

$$x \in \{0, 1\}^n, \tag{16c}$$

where $c_j \in \mathbb{R}$ and $a_j \in \{0, 1\}^m$, $j = 1, \dots, n$. Its Lagrangian with respect to the relaxation of the linear constraints (16b) has the form $L(x, u) := (1^m)^T u + \bar{c}^T x$, $u \in \mathbb{R}^m$, where we have defined the reduced cost vector $\bar{c} := c - A^T u$. Here, $c = (c_j)_{j=1}^n$ and $A = (a_1 a_2 \cdots a_n)$.

We define the Lagrangian dual problem to find

$$\theta^* := \text{maximum} \theta(u),$$

$$\text{subject to} \quad u \geq 0^m,$$

where

$$\theta(u) := (\mathbf{1}^m)^T u + \sum_{j=1}^n \underset{x_j \in \{0,1\}}{\text{minimum}} \bar{c}_j x_j, \quad u \geq 0^m,$$

is the dual function; the Lagrangian problem is of course solved such that

$$x_j(u) \begin{cases} = 1, & \text{if } \bar{c}_j < 0, \\ \in \{0, 1\}, & \text{if } \bar{c}_j = 0, \\ = 0, & \text{if } \bar{c}_j > 0. \end{cases} \quad (17)$$

In the following, we will report upon two experiments with new Lagrangian heuristics for the set covering problem. While these experiments were certainly not performed in order to establish the superiority of these new heuristics for solving large-scale set covering problems, the result encourages their use, especially when taking into account how easy it was to incorporate them into the well-established greedy strategy for set covering problems. Further developments of greedy heuristics along these lines, for this and other integer and combinatorial optimization problems, are left for future study.

Throughout, we have worked with the set covering problem `rail507`, for which we have the following data: $n = 63,009$, $m = 507$, and the best bounds reported in the literature are [172.1456, 174] (see [AMRT01] and [CFT99], respectively; the lower bound of 172.4 reported in [CNS98] is probably incorrect). Experiments—not reported on here—that we have performed on the similar set covering problems `rail516` and `rail582` corroborate the conclusions that are made below in Sections 4.4 and 4.5.

4.2 A generic primal greedy heuristic

A primal greedy heuristic is often a main component of a set covering algorithm. In our experiments, we will use several such algorithms, some of which are classic, and all of which can be written (at least essentially) as instances of the following generic primal greedy heuristic:

(Input) A primal vector $\bar{x} \in \{0, 1\}^n$ and a cost vector $p \in \mathbb{R}^n$.

(Output) A vector $\hat{x} \in \{0, 1\}^n$, feasible in (16).

(Starting phase) Given \bar{x} , delete all rows i in (16b) that are covered. Delete all variables x_j with $\bar{x}_j = 1$.

(Greedy insertion) Identify an undeleted variable x_τ which has the minimal value of p_j relative to the number (k_j) of uncovered, undeleted rows which it covers. Set $x_\tau := 1$. Delete all rows in (16b) which have been covered. Delete x_τ . If any uncovered rows in (16b) remain, then repeat this step; otherwise, let $\tilde{x} \in \{0, 1\}^n$ denote the feasible solution found.

(Greedy deletion in over-covered rows) Identify a variable x_τ with $\tilde{x}_\tau = 1$ which is present only in rows which are over-covered at \tilde{x} , and which has the maximal value of p_j relative to k_j . Set $\tilde{x}_\tau := 0$. If any such variable remains, then repeat this step; otherwise, let $\hat{x} \in \{0, 1\}^n$ denote the feasible solution found, and terminate.

We can identify several known instances of the above algorithm:

- (I) Let $\bar{x} := 0^n$ and $p := c$. This procedure is described by Chvátal [Chv79].
- (II) Let $\bar{x} := 0^n$ and $p := \bar{c}$, defined at some dual vector u (see above). This is essentially the heuristic PRIMAL5 of Balas and Ho [BaH80].
- (III) Let $\bar{x} := x(u)$, whose component $x_j(u)$ is given by (17), and let $p := c$. This heuristic is described by Beasley [Bea87, Bea93] and Wolsey [Wol98, Section 10.4].
- (IV) Let $\bar{x} := x(u)$ and $p := \bar{c}$. This is essentially the heuristic ERCGH of Balas and Carrera [BaC96].

The greedy selection criterion utilized in our description above is often of the form p_j/k_j , but can combine the entities p_j and k_j in different ways; such versions of the procedure (I) can be found in Balas and Ho [BaH80], and Vasko and Wilson [VaW84], where c_j/k_j is replaced by, among other choices, $c_j/\log_2 k_j$. (See [BaC96, Section 4] for an account of numerical experience with such heuristics.)

4.3 A simple dual algorithm

To find a good lower bound on θ^* , we apply *conditional subgradient optimization* ([LPS96]), using Polyak [Pol69] step lengths, that is, starting from a $u^0 \in \mathbb{R}_+^m$,

$$u^{t+1} := [u^t + \ell_t \gamma_+^t(u^t)]_+, \quad t = 0, 1, \dots,$$

where $[\cdot]_+$ denotes the Euclidean projection onto \mathbb{R}_+^m ,

$$(\gamma_+^t(u^t))_i := \begin{cases} 0, & \text{if } u_i^t = 0 \text{ and } a^i x(u^t) > 1, \\ 1 - a^i x(u^t), & \text{otherwise,} \end{cases} \quad i = 1, 2, \dots, m,$$

is a *projected subgradient* of θ with respect to \mathbb{R}_+^m at u^t , a^i being row i of the matrix A , and where the step length is

$$\ell_t := \nu_t \frac{UBD - \theta(u^t)}{\|\gamma_+^t(u^t)\|^2}, \quad t = 0, 1, \dots,$$

where the parameter $\nu_t := 1.5 \cdot 0.99^t$, $t = 0, 1, \dots$.

The upper bound UBD used in the step length formula was calculated a priori by applying the greedy heuristic (I) above. For the instance `rai1507`, the primal heuristic (I) produces $UBD = 209$. (The feasible solution found after the greedy insertion phase has the objective value 216.)

This subgradient algorithm is, in our experiments, terminated after a fixed number of iterations, a number which is altered among the experiments. The dual vector u chosen at termination was the final iteration point.

4.4 Experiment I

We have little information a priori about complementarity fulfillment and Lagrangian optimality at an optimal solution. To ensure that many possibilities are considered we will use an objective in the heuristic search that combines the Lagrangian function $[L(\cdot, u)]$ and complementarity fulfillment $[-u^T g(\cdot)]$, that is,

$$h(x) := \lambda[f(x) + u^T g(x)] + (1 - \lambda)[-u^T g(x)], \quad 0 \leq \lambda \leq 1. \quad (18)$$

In the set covering application, $h(x) = [\lambda \bar{c} + (1 - \lambda)A^T u]^T x$ holds. We obtain the original cost c by the choice $\lambda := 1/2$, while the Lagrangian reduced cost \bar{c} follows from the choice $\lambda := 1$. As remarked above, we will also consider other values. (Rewriting the expression in (18) as $\lambda[f(x) + \frac{2\lambda-1}{\lambda}u^T g(x)]$, we see that it can be interpreted as a Lagrangian if we choose $\lambda \geq 1/2$; in our experiments reported below, we found that the heuristic provided quite inferior solutions for $\lambda < 1/2$, and those solutions are not reported.)

Based on the generic primal heuristic of Section 4.2, we then define a set of heuristics that use a cost vector of the form $p := \lambda \bar{c} + (1 - \lambda)A^T u$. Note that all four instances (I)–(IV) described above use cost coefficients that are defined at the end-points of the interval for λ : (I) and (III) correspond to $\lambda := 1/2$, while (II) and (IV) correspond to $\lambda := 1$. In this first experiment, we define $\bar{x} := 0^n$, so the heuristics in this first test are all radical, according to our definition; the experiments in the next subsection also look at conservative heuristics, where $\bar{x} := x(u)$.

We ran two tests, the first with $t = 200$ as the final iteration, and the second with $t = 500$. For each of these two final values, we ran the above heuristic with values of λ in $[1/2, 1]$, with an increment of 0.005. The results are shown in Figure 3.

In Figure 3(a), the horizontal, dotted line is the value of the last lower bound found, which in this case was $\theta(u^{200}) = 159.53$. The three other lines, taken from the highest to the lowest, show the objective values of the feasible solutions obtained by the proposed Lagrangian heuristic, the values of δ , and the values of ε , as the value of λ ranges from $1/2$ to 1. The scale on the y -axis only applies to the original cost; for the latter two lines, the value zero corresponds to the dotted line. This line is, for Figure 3(b), at the level $\theta(u^{500}) = 170.84$. Both figures also illustrate that $f(x) = \theta(u) + \delta + \varepsilon$.

After 200 iterations, the quality of the lower bound is quite poor, and, running the primal heuristic from better dual solutions, we have found that better primal solutions are then always provided.

According to Figures 3(a) and (b), the value of ε clearly decreases with an increase in the value of λ , which is expected, since for $\lambda = 1$ the merit function used is the Lagrangian. The variation in δ (which always dominates in value over ε here) is less regular, except for larger values of λ when it increases rapidly, again as expected.




Figure 3: The Lagrangian heuristic for the set covering problem `rail507`. (a) $t = 200$; (b) $t = 500$.

The value of ε is very small for solutions that are of high quality, meaning that near-optimal solutions to this set covering problem violate complementarity to a large extent. It has indeed been observed that, for this class of problems, often several rows are over-covered in an optimal solution ([Tak01]).

The result of applying heuristic (I) is a feasible solution with cost 209; this corresponds to the height of the uppermost line at $\lambda = 1/2$ in both figures. The result of using heuristic (II) can be seen as the height of the uppermost line at $\lambda = 1$ in both figures (with the objective values 210 and 204, respectively). Clearly, both of these choices are inferior to using values of λ in the open interval $(1/2, 1)$. The best solutions are found for relatively large values of λ , as long as they are not very close to 1. The above observations lead us to advocate the use of heuristics based neither on the original cost, nor on the Lagrangian cost, but on a combination of the two of them, most importantly because this combination takes both Lagrangian optimality and complementarity violation into consideration.

The following experiment takes these observations as its starting point.

4.5 Experiment II

Based on the previous experiment, we chose to set $\lambda := 0.9$, and performed a second experiment. We ran three primal greedy heuristics at every iteration of the dual algorithm, starting from $t = 200$ and terminating at $t = 500$. For each of these, we recorded the objective value of the feasible solution \hat{x} obtained, and created histograms, as can be seen in Figure 4. The top one, (a), was obtained by using heuristic (III), that is, starting at the solution $\bar{x} := x(u^t)$ and using the original cost coefficients, $p := c$. This is a conservative heuristic. The middle figure (4b), was obtained by the instance of the generic heuristic where $\bar{x} := x(u^t)$ but where $p := \lambda \bar{c}^t + (1 - \lambda)A^T u^t$ with $\lambda = 0.9$. This is also a conservative heuristic, which, however, is based on a better merit function. (A preliminary experimental study indicated that the result of using this second heuristic—which was not evaluated in the previous section—was rather insensitive to the value of λ , as long as it was not chosen too close to either $1/2$ or 1 ; the value $\lambda = 0.9$ is a good compromise.) The bottom figure (4c) was created by the use of the primal heuristic, which also uses cost vector p , but which is radical because it takes $\bar{x} := 0^n$ as the starting point.

We observe from the figure a quite remarkable difference between the three heuristics; note in particular that the radical one consistently provides feasible solutions of rather high quality. For each of the three respective histograms, we have the following maximum, mean, and minimum objective values:

	(a)	(b)	(c)
maximum :	221	212	195
mean :	203.99	194.45	186.55
minimum :	192	182	182

Hence, the radical heuristic produces solutions the worst of which is nearly as good as the best outcome of the heuristic (III).

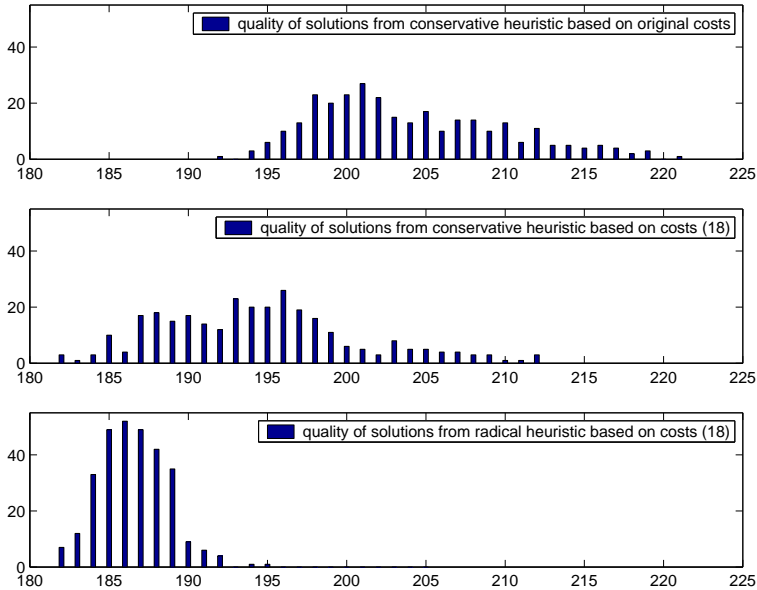


Figure 4: Quality of primal solutions obtained by three greedy heuristics. (a) $\bar{x} := x(u^t)$ and $p := c$ [heuristic (III)]; (b) $\bar{x} := x(u^t)$ and $p := 0.9\bar{c}^t + 0.1A^T u^t$; (c) $\bar{x} := 0^n$ and $p := 0.9\bar{c}^t + 0.1A^T u^t$.

We have also, in Figure 5, for each iteration plotted the moving average of the objective values over 30 iterations. (The value plotted for iteration 230 corresponds to the average of the iterations 201–230.) Figure 5 reveals that the radical heuristic provides relatively good primal solutions already at an early stage of the dual algorithm, and it clearly improves upon the other two. Further, the second of the conservative heuristics (which is also new) is, in turn, much better than the first, while being a very simple modification thereof. A general observation is that the primal solutions given by the conservative heuristics improve with the dual solution.

5 Applications to column generation and core problems

5.1 Column generation

The principle of column generation is most frequently used for attacking discrete optimization problems. (See, for example [Wol98, Chapter 11] and the recent survey in [Wil01].) Since, however, column generation is founded on linear programming duality, it is merely a continuous relaxation that can be solved by means of this principle. The results to be presented below introduce a degree of control over the integer programming quality of a column generation scheme.

Consider a discrete optimization problem with a feasible region that is defined by a finite Cartesian product set and a number of linear and coupling side constraints, that is, a problem of the form

$$f^* := \text{minimum} \sum_{j=1}^n c_j^T x_j, \quad (19a)$$

$$\text{subject to} \sum_{j=1}^n A_j x_j \geq b, \quad (19b)$$

$$x_j \in X_j, \quad j = 1, \dots, n, \quad (19c)$$

where the sets $X_j \subset \mathbb{R}^{n_j}$, $j = 1, \dots, n$, are finite, $c_j \in \mathbb{R}^{n_j}$ and $A_j \in \mathbb{R}^{m \times n_j}$, $j = 1, \dots, n$, and $b \in \mathbb{R}^m$. The problem is assumed to have a feasible solution.

In many applications that give rise to a model of this form, the sets X_j , $j = 1, \dots, n$, are described by linear constraints and integrality restrictions. The result to be presented below does not require this description to have the integrality property. Further, the objective function and the side constraints are




Figure 5: Moving averages of the solution quality for three greedy heuristics.

stated as being linear only in order to facilitate the presentation. (In the case of nonlinearities, additive separability over the Cartesian product is, however, required.)

For $j = 1, \dots, n$, let P_j denote the number of points in the set X_j , and denote these points by x_j^i , for $i = 1, \dots, P_j$. Problem (19) is then equivalent to the disaggregated master problem

$$f^* = \text{minimum} \sum_{j=1}^n \sum_{i=1}^{P_j} (c_j^T x_j^i) \lambda_j^i, \quad (20a)$$

$$\text{subject to} \sum_{j=1}^n \sum_{i=1}^{P_j} (A_j x_j^i) \lambda_j^i \geq b, \quad (20b)$$

$$\sum_{i=1}^{P_j} \lambda_j^i = 1, \quad j = 1, \dots, n, \quad (20c)$$

$$\lambda_j^i \in \{0, 1\}, \quad i = 1, \dots, P_j, \quad j = 1, \dots, n. \quad (20d)$$

Let $u \in \mathbb{R}_+^m$ be multipliers associated with the side constraints (19b), define the Lagrangian problem

$$\theta(u) := b^T u + \sum_{j=1}^n \theta_j(u), \quad (21)$$

with

$$\theta_j(u) := \text{minimum}_{x_j \in X_j} (c_j^T - u^T A_j) x_j, \quad j = 1, \dots, n,$$

and suppose that $\bar{u} \in \mathbb{R}_+^m$ is near-optimal in the Lagrangian dual problem to maximize the value of $\theta(u)$ over $u \in \mathbb{R}_+^m$.

Suppose further that $p_j \geq 1$ points in the respective sets X_j , $j = 1, \dots, n$, are available explicitly, let

$\delta \in \mathbb{R}_+$, and consider the following restricted master problem.

$$f_r^* := \text{minimum} \sum_{j=1}^n \sum_{i=1}^{p_j} (c_j^T x_j^i) \lambda_j^i, \quad (22a)$$

$$\text{subject to} \sum_{j=1}^n \sum_{i=1}^{p_j} (A_j x_j^i) \lambda_j^i \geq b, \quad (22b)$$

$$\sum_{j=1}^n \sum_{i=1}^{p_j} (\bar{u}^T A_j x_j^i) \lambda_j^i \leq \bar{u}^T b + \delta, \quad (22c)$$

$$\sum_{i=1}^{p_j} \lambda_j^i = 1, \quad j = 1, \dots, n, \quad (22d)$$

$$\lambda_j^i \in \{0, 1\}, \quad i = 1, \dots, p_j, \quad j = 1, \dots, n. \quad (22e)$$

The purpose of this nonstandard formulation of a restricted master problem, which has one side constraint, is that any feasible solution to it will satisfy near-complementarity [cf. (5c)].

This problem can be built up by, for example, enumerating (ranking) points in the product sets according to increasing objective values, or by applying column generation (to optimality or truncated) to the linear programming relaxation of the master problem (20). Another alternative is to apply subgradient optimization to the Lagrangian dual problem to maximize the function θ defined in (21) over \mathbb{R}_+^m , and, during the course of this scheme, accumulate optimal solutions to the relaxed problems (all of them or some only). (The use of subgradient optimization for accumulating columns to a linear programming restricted master problem is justified by, for example, [LaL97, Proposition 7]; see also [PeP97, LPS99, LPS03] for similar results for subgradient optimization applied to more general problems, [ShC96] for a related result in linear programming, and [Kiw95, FeK00] for primal recovery results using proximal dual subgradient methods.) Still another alternative for creating a restricted master problem is to enumerate the points in the product sets according to the reduced costs that are obtained within the column generation or Lagrangian relaxation approaches (cf. the discussion in Remark 9). [This enumeration may also filter out columns that can never be included in a feasible solution to the problem (20).] One might, of course, also consider combinations of all or some of these strategies. The result to be given below is valid, whichever principle has been used for constructing the restricted master problem (22).

REMARK 13 (observations). (a) If the Lagrangian relaxation (21) does not have the integrality property, then problem (20) is a strong formulation of problem (19), that is, in this case the linear programming relaxation of (20) is better than that of (19) in terms of the relative strength of their respective lower bounds. For further discussions about strong formulations, see, for example, [Wol98, Sections 1.7 and 10.2].

(b) We mention as a special example of the above the restricted master problem defined by Sweeney and Murphy [SwM79] in their decomposition method for integer programs; they build their restricted master problem by enumerating the vectors in the respective sets X_j according to a ranking based on Lagrangian reduced costs. [Their master problem is however the standard restriction of problem (20).]

(c) An interesting application of our column generation methodology is its combination with the branch-and-price algorithm of Savelsbergh [Sav97] (see also [BJNSV98]) for the generalized assignment problem wherein knapsack solutions (columns) are generated inside a tree search. In Savelsbergh's algorithm, the initial node comprises the solution of the linear programming relaxation of the master problem through (traditional) column generation. In our proposed combination, one creates a larger restricted master problem at the root node by also enumerating knapsack solutions, through ranking them with respect to their reduced costs. Our alternative approach yields a larger initial master problem but reduces, or completely eliminates, the need for a further column generation in the branch-and-bound search. \square

In order to state the main result, we define the (Lagrangian) reduced costs

$$\bar{c}_j^i := (c_j^T - \bar{u}^T A_j) x_j^i - \theta_j(\bar{u}), \quad i = 1, \dots, p_j, \quad j = 1, \dots, n, \quad (23)$$

for the variables in the restricted master problem. (Note that $\bar{c}_j^i \geq 0$ always holds.)

THEOREM 14 (quality of restricted master problem). *Let $\bar{u} \in \mathbb{R}_+^m$ and $\delta \in \mathbb{R}_+$. Suppose that the restricted master problem (22) has a feasible solution. Then,*

$$f_r^* \leq \theta(\bar{u}) + \varepsilon + \delta,$$

where

$$\varepsilon := \sum_{j=1}^n \max_{i=1, \dots, p_j} \bar{c}_j^i. \quad (24)$$

PROOF. We utilize Proposition 5(a), as follows. Let $x_j = x_j^{i(j)}$, where $i(j) \in \{1, \dots, p_j\}$, for $j = 1, \dots, n$, be a feasible solution to problem (19) that corresponds to an optimal solution to (22). Denote this solution by \bar{x} . By using (23) and (21) we then have that

$$\begin{aligned} L(\bar{x}, \bar{u}) &:= b^T \bar{u} + \sum_{j=1}^n (c_j^T - \bar{u}^T A_j) x_j^{i(j)} = b^T \bar{u} + \sum_{j=1}^n \left(\theta_j(\bar{u}) + \bar{c}_j^{i(j)} \right) \\ &\leq \theta(\bar{u}) + \sum_{j=1}^n \max_{i=1, \dots, p_j} \bar{c}_j^i = \theta(\bar{u}) + \varepsilon. \end{aligned}$$

Hence, (10a) holds. Further, (10b) holds by assumption. Finally, with $g(x) := b - \sum_{j=1}^n A_j x_j$, (22c) gives that $\bar{u}^T g(\bar{x}) \geq -\delta$ holds, whence (10c) follows. Hence, the pair (\bar{x}, \bar{u}) satisfies (10), and the conclusion follows. \square

In the case for which inequality constraints (19b) are replaced by equalities, the restricted master problem is modified accordingly, (22c) is not present, and $\delta = 0$.

The theorem immediately implies that

$$f_r^* - f^* \leq \left(\sum_{j=1}^n \max_{i=1, \dots, p_j} \bar{c}_j^i \right) + \delta$$

holds.

The result below provides a (limited) possibility to assess the quality of the restricted master problem. The result can also be used as a guide to the adjustment of the restriction; especially, the result describes a property of the restricted master problem such that its feasible set contains an optimal solution to (19). The proof of the result is straightforward, and is therefore omitted.

PROPOSITION 15 (variable fixing and optimality). *Let $\bar{u} \in \mathbb{R}_+^m$. Suppose that the restricted master problem (22) has a feasible solution, and let \bar{x} be the feasible solution to (19) that corresponds to an optimal solution to (22). Suppose further that we know an upper bound $\bar{f} \geq f^*$ (for example, $\bar{f} = f_r^*$).*

- (a) *If a $j \in \{1, \dots, n\}$ is such that $\bar{c}_j^i > \bar{f} - \theta(\bar{u})$ holds for $i = p_j + 1, \dots, P_j$, then $x_j^* = \bar{x}_j$ in every optimal solution x^* to (19).*
- (b) *If $\delta \geq \bar{f} - \theta(\bar{u})$ and if $\bar{c}_j^i \geq \bar{f} - \theta(\bar{u})$ holds for $i = p_j + 1, \dots, P_j$ and for every $j \in \{1, \dots, n\}$, then $x^* = \bar{x}$ is an optimal solution to (19). \square*

REMARK 16 (observations and future research). (a) The obvious way to tighten the upper bound in the theorem is to delete from the restricted master problem the columns with maximal reduced costs (23), and to decrease the value of δ ; such further restrictions of (22) might, however, cause the problem to become infeasible.

(b) Whenever a feasible solution to (19) is known, it can be utilized to define vectors in X_j and a value of δ such that the initial restricted master problem (22) has a feasible solution. Note also that the columns corresponding to the solutions to the relaxed problems

$$\theta_j(\bar{u}) := \min_{x_j \in X_j} (c_j^T - \bar{u}^T A_j) x_j, \quad j = 1, \dots, n,$$

need not be present in (22) [although it is, of course, quite natural that they are].

(c) The global optimality conditions (5) clearly indicate that the traditional means of generating columns in column generation algorithms for an integer program with inequality constraints—that of finding Lagrangian optimal solutions—is too limited: complementarity (near-)fulfillment is just as an

important goal. These two goals can of course be taken into account in different ways, leading to different column generation algorithms. One example column generation problem is a Lagrangian-type problem whose objective function is a convex combination of the two goals; we used it in our set covering heuristic.

An interesting avenue for future research is the development of column generation algorithms for the integer program (20), that are elicited by the structure of the side-constrained restricted master problem (22). Consider, for example, a side constrained extension to the master problem (20), that is, problem (22) with $p_j = P_j$ for all j . Suppose further that we Lagrangian relax side constraint (22c) with a fixed, positive multiplier ($\eta > 0$). The resulting master problem then has an objective function that combines the Lagrangian function in (21) and a complementarity term; this combination is similar to the function h which was utilized in the development of set covering heuristics in the previous section. When applying column generation to this master problem, this combination explicitly incorporates near-complementarity fulfillment in the objective function of the column generation problem. Our conjecture is that the columns thus constructed should be better from the point of view of solving the original integer program than those that are generated in a column generation algorithm based on the continuous relaxation of (20). With reference to our experiments on the set-covering problem, it is probably beneficial to allow the value of the above-mentioned multiplier (η) to vary somewhat in the column generation for a given value of \bar{u} (whence the multiplier assumes the same role as λ in that application), as it would allow us to generate columns that are near-optimal in the Lagrangian problem and/or near-complementary. If the value of η is chosen such that the column generation subproblem correctly and optimally weighs between Lagrangian near-optimality and near-complementarity, then the side constraint (22c) can be removed from the restricted master problem (22).

The resulting column generation scheme is based on the standard restricted master problem, but has a column generation subproblem whose objective function combines the classic Lagrangian term with a complementarity term.

There are of course other means by which both near-optimality and near-complementarity can be taken into account in column generation, for example via ranking and filtering, or even via explicit constraints in the column generation problem. \square

We end with a corollary to Proposition 15, which stems from the application of column generation to the LP relaxation of (20). Suppose we have solved to (near-)optimality the linear program that is the LP relaxation of the restricted master problem (22), and that we have access to a primal basic feasible solution to this linear program with objective value f_{RMP} , and a complementary dual solution $(\hat{u}, \hat{v}, \hat{w})$. Let the best column obtained in the column-generation phase for the LP relaxation of problem (20), based on the dual variable vector \hat{u} , have the reduced cost $\bar{c}_j^{p_j+1}$ ($j = 1, \dots, n$). (If the LP relaxation of the restricted master problem (22) was solved to optimality and its solution is not optimal in the LP relaxation of (20), then the newly generated column had not been generated previously.) Consider then the quality of the next restricted master problem:

COROLLARY 17 (solution quality in column generation). *In the current setting, we have the estimate*

$$f_{\text{RMP}} + \sum_{j=1}^n \bar{c}_j^{p_j+1} \leq f^* \leq f_r^* \leq f_{\text{RMP}} + \sum_{j=1}^n \left(\bar{c}_j^{p_j+1} + \max_{i=1, \dots, p_j} \bar{c}_j^i \right) + \delta$$

of the optimal value. \square

5.2 Core problems

The formulation and solution of core problems is a means of utilizing Lagrangian duality to find primal feasible and near-optimal solutions which is more sophisticated than simpler, manipulative heuristics such as those presented in Section 4. As such, core problems define Lagrangian heuristics which may or may not be conservative, depending on the principle used for constructing the core problem, and on the size of the latter. Core problems lie at the heart of the set-covering heuristics in [CNS98, CFT99], and efficient “core algorithms” have also been developed for binary knapsack problems (e.g., [BaZ80, MaT90, Pis95]). These core problems are based primarily on the (linear programming-)reduced costs of the variables, and so they can be said to focus on Lagrangian near-optimality as opposed to also incorporating complementarity near-fulfillment. Our analysis of core problems suggests that complementarity near-fulfillment can, and should, be introduced in core problems.

Consider the binary problem

$$f^* := \text{minimum} \sum_{j=1}^n c_j x_j, \quad (25a)$$

$$\text{subject to} \sum_{j=1}^n a_j x_j \geq b, \quad (25b)$$

$$\sum_{j=1}^n d_j x_j \geq e, \quad (25c)$$

$$x \in \{0, 1\}^n, \quad (25d)$$

where b, a_j ($j = 1, \dots, n$) $\in \mathbb{R}^m$ and e, d_j ($j = 1, \dots, n$) $\in \mathbb{R}^r$. Suppose that the problem is feasible. We propose solving this problem by means of a Lagrangian relaxation of the (complicating) constraints (25b), the multipliers being $u \in \mathbb{R}_+^m$. We assume that the resulting Lagrangian problem

$$\theta(u) := b^T u + \text{minimum}_{\substack{\sum_{j=1}^n d_j x_j \geq e \\ x_j \in \{0, 1\}}} \sum_{j=1}^n (c_j - u^T a_j) x_j, \quad u \geq 0^m, \quad (26)$$

has the integrality property, so that each constraint $x_j \in \{0, 1\}$ can be replaced by $0 \leq x_j \leq 1$, for all $j = 1, \dots, n$, without affecting the value of the dual function. We denote the corresponding linear programming dual multipliers for constraints (25c) by $v \in \mathbb{R}_+^r$. At a near-optimal $\bar{u} \in \mathbb{R}_+^m$ in the Lagrangian dual problem to maximize the value of $\theta(u)$ over $u \in \mathbb{R}_+^m$, let $(x(\bar{u}), \bar{v})$ be the optimal primal-dual solution to the Lagrangian problem (26).

The construction of the core problem is based on the idea of predicting, that is, guessing the optimal values of some of the x_j variables in problem (25); the value of $x(\bar{u})$, or the reduced cost vector, is most often used for this purpose. We denote by \mathcal{J}_0 (\mathcal{J}_1) those indices $j \in \mathcal{J} := \{1, \dots, n\}$ for which the *prediction* is that $x_j^* = 0$ (respectively, $x_j^* = 1$). A core problem is a restriction of the original problem (25) wherein the variables in $\mathcal{J}_0 \cup \mathcal{J}_1$ are fixed to their predicted values, and the remaining variables, $\mathcal{J}_f := \mathcal{J} \setminus (\mathcal{J}_0 \cup \mathcal{J}_1)$, are free. We assume that Lagrangian problem (26) has the integrality property, in order to be able to utilize reduced costs in the ranking of the variables when deciding on the predictions.

Let $\Delta_1 \in \mathbb{R}_+^m$, $\Delta_2 \in \mathbb{R}_+^r$. Our core problem is

$$f_c^* := \sum_{j \in \mathcal{J}_1} c_j + \text{minimum} \sum_{j \in \mathcal{J}_f} c_j x_j, \quad (27a)$$

$$\text{subject to} \quad b - \sum_{j \in \mathcal{J}_1} a_j + \Delta_1 \geq \sum_{j \in \mathcal{J}_f} a_j x_j \geq b - \sum_{j \in \mathcal{J}_1} a_j, \quad (27b)$$

$$e - \sum_{j \in \mathcal{J}_1} d_j + \Delta_2 \geq \sum_{j \in \mathcal{J}_f} d_j x_j \geq e - \sum_{j \in \mathcal{J}_1} d_j, \quad (27c)$$

$$x_f \in \{0, 1\}, \quad j \in \mathcal{J}_f. \quad (27d)$$

As was the case with the restricted master problem (22), the purpose of this core problem construction is that feasible solutions are near-complementary. It is to be noted that the means by which we here enforce near-complementarity is slightly different from the one in (22); a constraint like (22c) is avoided here because it would destroy any favourable structure inherent in constraints (27b)–(27c).

Although, in principle, the subsets \mathcal{J}_0 and \mathcal{J}_1 can be defined quite arbitrarily, it is natural to choose them such that

$$\mathcal{J}_0 \subseteq \{j \in \mathcal{J} \mid \bar{c}_j > 0\}, \quad (28a)$$

$$\mathcal{J}_1 \subseteq \{j \in \mathcal{J} \mid \bar{c}_j < 0\}, \quad (28b)$$

where

$$\bar{c}_j := c_j - \bar{u}^T a_j - \bar{v}^T d_j, \quad j \in \mathcal{J} \quad (29)$$

define the (linear programming-)reduced cost vector in the Lagrangian problem. (An obvious means for creating core problems according to principle (28) is to compare the values of the reduced costs \bar{c}_j with some non-zero threshold value.)

THEOREM 18 (quality of core problem). *Let $\bar{u} \in \mathbb{R}_+^m$, and let $(x(\bar{u}), \bar{v})$ solve the Lagrangian problem (26). Suppose that the prediction satisfies (28) and that the core problem (27) has a feasible solution. Then,*

$$f_c^* \leq \theta(\bar{u}) + \varepsilon + \delta,$$

where

$$\begin{aligned} \varepsilon &:= \sum_{j \in \mathcal{J}_f} |\bar{c}_j| + \bar{v}^\top \Delta_2, \\ \delta &:= \bar{u}^\top \Delta_1. \end{aligned}$$

PROOF. The proof utilizes Proposition 5(a), as follows.

Let $\bar{x} := (\bar{x}_{\mathcal{J}_0}, \bar{x}_{\mathcal{J}_1}, \bar{x}_{\mathcal{J}_f})$ be the primal vector corresponding to the predictions and the optimal solution to the core problem (27).

To establish that (10a) holds, we use

$$\bar{v}^\top \left(\sum_{j \in \mathcal{J}} d_j x_j(\bar{u}) - e \right) = 0, \quad \text{and} \quad \bar{v}^\top \left(e - \sum_{j \in \mathcal{J}} d_j \bar{x}_j + \Delta_2 \right) \geq 0,$$

holds, where the equality follows from the complementarity conditions in the linear programming problem equivalent to (26), and the inequality follows from the feasibility of \bar{x} in (27); in particular,

$$\bar{v}^\top \sum_{j \in \mathcal{J}} d_j [x_j(\bar{u}) - \bar{x}_j] + \bar{v}^\top \Delta_2 \geq 0. \quad (30)$$

We therefore have that

$$\begin{aligned} L(\bar{x}, \bar{u}) &:= b^\top \bar{u} + \sum_{j \in \mathcal{J}} (c_j - \bar{u}^\top a_j) \bar{x}_j = \theta(\bar{u}) + \sum_{j \in \mathcal{J}} (c_j - \bar{u}^\top a_j) [\bar{x}_j - x_j(\bar{u})] \\ &\leq \theta(\bar{u}) + \sum_{j \in \mathcal{J}} \bar{c}_j [\bar{x}_j - x_j(\bar{u})] + \bar{v}^\top \Delta_2 \quad [(29) \text{ and } (30)] \\ &\leq \theta(\bar{u}) + \sum_{j \in \mathcal{J}_f} |\bar{c}_j| + \bar{v}^\top \Delta_2 \quad [(28)] \\ &= \theta(\bar{u}) + \varepsilon. \end{aligned}$$

That (10b) holds follows by assumption.

Finally, from (27b), we obtain that, with $g(x) := b - \sum_{j \in \mathcal{J}} a_j x_j$, the relation $g(\bar{x}) \geq -\Delta_1$ holds, and so we obtain, because $\bar{u} \in \mathbb{R}_+^m$, that $\bar{u}^\top g(\bar{x}) \geq -\bar{u}^\top \Delta_1 = -\delta$, hence (10c) follows. Thus, the pair (\bar{x}, \bar{u}) satisfies (10).

Proposition 5(a) then yields the desired result. \square

The theorem immediately implies that

$$f_c^* - f^* \leq \sum_{j \in \mathcal{J}_f} |\bar{c}_j| + \bar{u}^\top \Delta_1 + \bar{v}^\top \Delta_2.$$

Note that if the core problem is constructed from an optimal dual solution (which can be found by solving the continuous relaxation of the original problem, due to the integrality property) and the values of Δ_1 and Δ_2 are taken to be at least as large as the LP optimal slacks, then the continuous relaxation of the core problem is feasible.

The following result corresponds to Proposition 15.

PROPOSITION 19 (variable fixing and optimality). *Let $\bar{u} \in \mathbb{R}_+^m$. Suppose that the prediction satisfies (28), that the core problem (27) has a feasible solution, and that \bar{x} is an optimal solution to it. Suppose further that we know an upper bound $\bar{f} \geq f^*$ (for example, $\bar{f} = f_c^*$).*

- (a) If a $j \in \mathcal{J}_0 \cup \mathcal{J}_1$ is such that $|\bar{c}_j| > \bar{f} - \theta(\bar{u})$, then $x_j^* = \bar{x}_j$ in every optimal solution x^* to (25).
(b) If

$$(\Delta_1)_i \begin{cases} \geq (\bar{f} - \theta(\bar{u}))/\bar{u}_i, & \text{if } \bar{u}_i > 0, \\ = \infty, & \text{otherwise,} \end{cases} \quad i = 1, \dots, m,$$

$$(\Delta_2)_i \begin{cases} \geq (\bar{f} - \theta(\bar{u}))/\bar{v}_i, & \text{if } \bar{v}_i > 0, \\ = \infty, & \text{otherwise,} \end{cases} \quad i = 1, \dots, r,$$

and if $|\bar{c}_j| \geq \bar{f} - \theta(\bar{u})$ holds for every $j \in \mathcal{J}_0 \cup \mathcal{J}_1$, then $x^* = \bar{x}$ is an optimal solution to (25). \square

An implication of this result is the well-known fact that core problems should comprise variables having small reduced costs.

As an ending note, we provide an important comment on the construction of the master and core problems in these last two subsections. [This comment is made relative to core problems, but the same arguments apply for the side-constrained master problem (22).] If the value of each of the elements of Δ_1 and Δ_2 increases, then the value of f_c^* reduces, since the feasible set of the problem (27) then would be enlarged; therefore, one might ask, “what is the purpose of constraints (27b) and (27c)?” The answer is that there is a trade-off between obtaining good objective values with a core problem and the complexity of solving it; constraints (27b) and (27c) make the problem easier to solve by restricting the feasible set in comparison to when (25b) and (25c) are present. Further, the constraints used in problem (27), and which serve to control the violation of complementarity, explicitly incorporate a measure that we have shown previously to be of utmost importance in forcing Lagrangian heuristics to approach an optimal solution, and a term which heretofore has not been incorporated in core problems, because only the Lagrangian optimality-based term defined by the vector of reduced costs \bar{c} is normally used.

Acknowledgements

Dr. Laura Wynter provided valuable comments on an early draft. Lic. Phil. Tuomo Takkula pointed out the relations between our Proposition 5(b) and Corollary 7 to [NeW88, Corollary II.3.6.9 and Theorem II.3.6.7]. Discussions with Lic. Phil. Myrna Palmgren inspired the study of relations to the work in [HaZ80] which lead to Remark 9. Lic. Phil. Andreas Westerlund drew our attention to the paper [SwM79].

The first author was partially financed by grant # 621-2001-2716 from the Swedish Science Foundation.

We thank the two anonymous reviewers and the associate editor for remarks that resulted in an improved presentation, and also the area editor Hanif Sherali for his enthusiasm and encouragement.

We dedicate this article to the Swedish scientist Leif H. Appelgren, for his pioneering work in [App69, App71] on column generation algorithms for integer and combinatorial optimization.

References

- [AEOP02] R. K. Ahuja, O. Ergun, J. B. Orlin, and A. P. Punnen. A survey of very large scale neighborhood search techniques. *Discrete Applied Mathematics*, 123 (2002), pp. 75–102.
- [AMRT01] E. D. Andersen, J. E. Mitchell, C. Roos, and T. Terlaky. A homogenized cutting plane method to solve the convex feasibility problem. In: *Optimization Methods and Applications*, X.-Q. Yang, K.L. Teo, and L. Caccetta (eds.), Kluwer Academic Publishers, Dordrecht, The Netherlands, 2001, pp. 167–190.
- [App69] L. H. Appelgren. A column generation algorithm for a ship scheduling problem. *Transportation Science*, 3 (1969), pp. 53–68.
- [App71] L. H. Appelgren. Integer programming methods for a vessel scheduling problem. *Transportation Science*, 5 (1971), pp. 64–78.
- [BaC96] E. Balas and M. C. Carrera. A dynamic subgradient-based branch-and-bound procedure for set covering. *Operations Research*, 44 (1996), pp. 875–890.

- [BaH80] E. Balas and A. Ho. An algorithm for set covering problems. *Mathematical Programming Study*, 12 (1980), pp. 37–60.
- [BaZ80] E. Balas and E. Zemel. An algorithm for large zero-one knapsack problems. *Operations Research*, 28 (1980), pp. 1130–1154.
- [BJNSV98] C. Barnhart, E. L. Johnson, G. L. Nemhauser, M. W. P. Savelsbergh, and P. H. Vance. Branch-and-price: Column generation for huge integer programs. *Operations Research* 46 (1998), pp. 316–329.
- [BSS93] M. S. Bazaraa, H. D. Sherali, and C. M. Shetty. *Nonlinear Programming: Theory and Algorithms*. John Wiley & Sons, New York, NY, second edition, 1993.
- [Bea87] J. E. Beasley. An algorithm for set covering problems. *European Journal of Operational Research*, 31 (1987), pp. 85–93.
- [Bea93] J. E. Beasley. Lagrangian relaxation. Chapter 6 in *Modern Heuristic Techniques for Combinatorial Problems*, C.R. Reeves (ed.), Blackwell Scientific Publications, Oxford, U.K., 1993.
- [Ber99] D. P. Bertsekas. *Nonlinear Programming*. Second edition. Athena Scientific, Belmont, MA, 1999.
- [BNO03] D. P. Bertsekas, A. Nedić, and A. E. Ozdaglar. *Convex Analysis and Optimization*. Athena Scientific, Belmont, MA, 2003.
- [BGB81] G. G. Brown, A. Geoffrion, and G. Bradley. Seasonal production and sales planning with limited shared tooling at the key operation. *Management Science*, 27 (1981), pp. 247–259.
- [CaM91] G. M. Campbell and V. A. Mabert. Cyclical schedules for capacitated lot sizing with dynamic demands. *Management Science*, 17 (1991), pp. 409–427.
- [CFT99] A. Caprara, M. Fischetti, and P. Toth. A heuristic method for the set covering problem. *Operations Research*, 47 (1999), pp. 730–743.
- [CFT02] A. Caprara, M. Fischetti, and P. Toth. Modeling and solving the train timetabling problem. *Operations Research*, 50 (2002), pp. 851–861.
- [CaW03] W. M. Carlyle and R. K. Wood. Lagrangian relaxation and enumeration for solving constrained shortest-path problems. In: *Proceedings of the 38th Annual ORSNZ Conference*, University of Waikato, Hamilton, New Zealand, 21–22 November, 2003.
- [CNS98] S. Ceria, P. Nobile, and A. Sassano. A Lagrangian-based heuristic for large-scale set covering problems. *Mathematical Programming*, 81 (1998), pp. 215–228.
- [Chv79] V. Chvátal. A greedy heuristic for the set-covering problem, *Mathematics of Operations Research*, 4 (1979), pp. 233–235.
- [Erl78] D. Erlenkotter. A dual-based procedure for uncapacitated facility location. *Operations Research*, 26 (1978), pp. 992–1009.
- [Eve63] H. Everett, III. Generalized Lagrange multiplier method for solving problem of optimal allocation of resources. *Operations Research*, 11 (1963), pp. 399–518.
- [FeK00] S. Feltenmark and K. C. Kiwiel. Dual applications of proximal bundle methods, including Lagrangian relaxation of nonconvex problems. *SIAM Journal on Optimization*, 14 (2000), pp. 697–721.
- [Fis81] M. L. Fisher. The Lagrangian relaxation method for solving integer programming problems. *Management Science*, 27 (1981), pp. 1–18.
- [Fis85] M. L. Fisher. An applications oriented guide to Lagrangian relaxation. *Interfaces*, 15 (1985), pp. 10–21.
- [Fis94] M. L. Fisher. Optimal solution of vehicle-routing problems using minimum K-trees. *Operations Research*, 42 (1994), pp. 626–642.
- [Geo74] A. M. Geoffrion. Lagrangian relaxation for integer programming. *Mathematical Programming Study*, 2 (1974), pp. 82–114.
- [HaZ80] G. Y. Handler and I. Zang. A dual algorithm for the constrained shortest path problem. *Networks*, 10 (1980), pp. 293–310.

- [HeK70] M. Held and R. M. Karp. The traveling salesman problem and minimum spanning trees: Part I. *Operations Research*, 18 (1970), pp. 1138–62.
- [HeK71] M. Held and R. M. Karp. The traveling-salesman problem and minimum spanning trees: Part II. *Mathematical Programming*, 1 (1971), pp. 6–25.
- [HWC74] M. Held, P. Wolfe and H. P. Crowder. Validation of subgradient optimization. *Mathematical Programming*, 16 (1974), pp. 62–88.
- [JLV90] Ö. Jonsson, T. Larsson, and P. Värbrand. A Lagrangean relaxation scheme for a scheduling problem. *Asia-Pacific Journal of Operations Research*, 7 (1990), pp. 155–162.
- [Kiw95] K. C. KIWIEL. Proximal level bundle methods for convex nondifferentiable optimization, saddle-point problems and variational inequalities. *Mathematical Programming*, 69 (1995), pp. 89–109.
- [LaL97] T. Larsson and Z.-W. Liu. A Lagrangean relaxation scheme for structured linear programs with application to multicommodity network flows. *Optimization*, 40 (1997), pp. 247–284.
- [LPS96] T. Larsson, M. Patriksson, and A.-B. Strömberg. Conditional subgradient optimization—theory and applications. *European Journal of Operational Research*, 88 (1996), pp. 382–403.
- [LPS99] T. Larsson, M. Patriksson, and A.-B. Strömberg. Ergodic, primal convergence in dual subgradient schemes for convex programming. *Mathematical Programming*, 86 (1999), pp. 283–312.
- [LPS03] T. Larsson, M. Patriksson, and A.-B. Strömberg. On the convergence of conditional ε -subgradient methods for convex programs and convex–concave saddle-point problems. *European Journal of Operational Research*, 151 (2003), pp. 461–473.
- [Las70] L. S. Lasdon. *Optimization Theory for Large Systems*. Macmillan, New York, NY, 1970.
- [MSW76] T. L. Magnanti, J. F. Shapiro, and M. H. Wagner. Generalized linear programming solves the dual. *Management Science*, 21 (1976), pp. 1195–1203.
- [Mar93] J. Marklund. A study of Lagrangian heuristics for convex network flow problems. Master’s thesis, Department of Mathematics, Linköping Institute of Technology, Linköping, Sweden, Linköping, Sweden, 1993.
- [MaT90] S. Martello and P. Toth. *Knapsack Problems: Algorithms and Computer Implementations*. Wiley, Chichester, UK, 1990.
- [MuC79] J. Mulvey and H. Crowder. Cluster analysis: An application of Lagrangian relaxation. *Management Science*, 25 (1979), pp. 329–340.
- [NeW88] G. L. Nemhauser and L. A. Wolsey. *Integer and Combinatorial Optimization*. John Wiley & Sons, New York, NY, 1988.
- [Pat94] M. Patriksson. *The Traffic Assignment Problem—Models and Methods*. Topics in Transportation. VSP BV, Utrecht, The Netherlands, 1994.
- [PeP97] J. Petersson and M. Patriksson. Topology optimization of sheets in contact by a subgradient method. *International Journal of Numerical Methods in Engineering*, 40 (1997), pp. 1295–1321.
- [Pis95] D. Pisinger. An expanding-core algorithm for the exact 0-1 knapsack problem. *European Journal of Operational Research*, 87 (1995), pp. 175–187.
- [Pol69] B. T. Polyak. Minimization of unsmooth functionals. *USSR Computational Mathematics and Mathematical Physics*, 9 (1969), pp. 14–29.
- [Sav97] M. W. P. Savelsbergh. A Branch-and-price algorithm for the generalized assignment problem. *Operations Research*, 45 (1997), pp. 831–841.
- [Sha79a] J. F. Shapiro. *Mathematical Programming: Structures and Algorithms*. John Wiley & Sons, New York, NY, 1979.
- [Sha79b] J. F. Shapiro. A survey of Lagrangian techniques for discrete optimization. *Annals of Discrete Mathematics*, 5 (1979), pp. 113–138.
- [ShC96] H. D. Sherali and G. Choi. Recovery of primal solutions when using subgradient optimization methods to solve Lagrangian duals of linear programs, *Operations Research Letters*, 19 (1996), pp. 105–113.

- [SNH83] J.-J. Strodiot, V. H. Nguyen, and N. Heukemes. ε -optimal solutions in nondifferentiable convex programming and some related questions. *Mathematical Programming*, 25 (1983), pp. 307–328.
- [SwM79] D. J. Sweeney and R. A. Murphy. A method of decomposition for integer programs. *Operations Research*, 27 (1979). pp. 1128–1141.
- [Tak01] T. Takkula. Private communication, 2001.
- [VaW84] F. J. Vasko and G. R. Wilson. An efficient heuristic for large set covering problems. *Naval Research Logistics*, 31 (1984), pp 163–171.
- [Ven91] J. A. Ventura. Computational development of a Lagrangian dual approach for quadratic networks. *Networks*, 21 (1991), pp. 469–485.
- [Wil01] W. E. Wilhelm. A technical review of column generation in integer programming. *Optimization and Engineering*, 2 (2001), pp. 159–200.
- [Wol98] L. A. Wolsey. *Integer Programming*. John Wiley & Sons, New York, 1998.