

Comparison of Bundle and Classical Column Generation

O. Briant, C. Lemaréchal, Ph. Meurdesoif, S. Michel, N. Perrot, F.
Vanderbeck

► **To cite this version:**

O. Briant, C. Lemaréchal, Ph. Meurdesoif, S. Michel, N. Perrot, et al.. Comparison of Bundle and Classical Column Generation. [Research Report] RR-5453, INRIA. 2005, pp.31. inria-00070554

HAL Id: inria-00070554

<https://hal.inria.fr/inria-00070554>

Submitted on 19 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Comparison of Bundle and Classical Column Generation

O. Briant — C. Lemaréchal — Ph. Meurdesoif — S. Michel — N. Perrot — F. Vanderbeck

N° 5453

Janvier 2005

THÈME 4



*Rapport
de recherche*

Comparison of Bundle and Classical Column Generation *

O. Briant[†], C. Lemaréchal[‡], Ph. Meurdesoif, S. Michel, N. Perrot, F. Vanderbeck[§]

Thème 4 — Simulation et optimisation
de systèmes complexes
Projet Bipop

Rapport de recherche n° 5453 — Janvier 2005 — 31 pages

Abstract: When a column generation approach is applied to decomposable mixed integer programming problems, it is standard to formulate and solve the master problem as a linear program. Seen in the dual space, this results in the algorithm known in the nonlinear programming community as the cutting-plane algorithm of Kelley and Cheney-Goldstein. However, more stable methods with better theoretical convergence rates are known and have been used as alternatives to this standard. One of them is the bundle method; our aim is to illustrate its differences with Kelley's method. In the process we review alternative stabilization techniques used in column generation, comparing them from both primal and dual points of view. Numerical comparisons are presented for five applications: cutting stock (which includes bin packing), vertex coloring, capacitated vehicle routing, multi-item lot sizing, and traveling salesman.

Key-words: Lagrangian duality, Dantzig-Wolfe decomposition, stabilized column generation, cutting-plane algorithms, nonsmooth convex optimization

* This research has been supported by Inria New Investigation Grant "Convex Optimization and Dantzig-Wolfe Decomposition"

[†] MAB, Univ. Bordeaux 1, 351 cours de la Libération, 33405 Talence

[‡] Inria Rhône-Alpes, 655 avenue de l'Europe, Montbonnot, 38334 Saint Ismiers

[§] MAB, Univ. Bordeaux 1, 351 cours de la Libération, 33405 Talence

Comparaison des Méthodes de Faisceaux avec la Génération de Colonnes Classique

Résumé : Lorsqu'une approche par génération de colonnes est appliquée à un problème décomposable en variables mixtes, le maître problème est traditionnellement formulé et résolu comme un programme linéaire. Vu dans l'espace dual, ceci donne l'algorithme connu en optimisation non linéaire sous le nom de plans sécants de Kelley, Cheney-Goldstein. Pourtant, des méthodes plus stables, et ayant un meilleur taux de convergence, sont connues et ont été utilisées à la place de cet usage courant. L'une d'entre elles est la méthode de faisceaux; notre but est d'illustrer ses différences avec la méthode de Kelley. Par la même occasion, nous passons en revue d'autres techniques de stabilisation utilisées en génération de colonnes, les comparant des points vue dual et primal. Nous présentons des comparaisons numériques sur cinq applications: découpe industrielle (qui inclut le bin packing), coloration de graphe, routage de véhicules avec capacité limitée, série économique multi-lots, et voyageur de commerce.

Mots-clés : Dualité lagrangienne, décomposition de Dantzig-Wolfe, génération de colonnes stabilisée, algorithmes de plans sécants, optimisation convexe non différentiable

1 Algorithms for column generation

This paper deals with optimization problems of the form

$$\min cx, \quad Ax \geq b \in \mathbb{R}^m, \quad x \in X := \{x^i : i \in I\} \subset \mathbb{R}^n, \quad (1)$$

where the index set I is assumed finite (think of X as being a bounded integer polyhedron). An equivalent formulation is Dantzig-Wolfe's (integer) *master problem*

$$\min \sum_{i \in I} (cx^i) \lambda_i, \quad \sum_{i \in I} (Ax^i) \lambda_i \geq b, \quad \sum_{i \in I} \lambda_i = 1, \quad \lambda_i \in \{0, 1\}, \quad i \in I,$$

in which the variable vector is $\lambda \in \{0, 1\}^{|I|}$.

When solving such problems, it is standard to make use of a lower bound obtained by *relaxation*. Denoting by $\text{conv}(X)$ the convex hull of the set X , the present paper is rather devoted to solving

$$\min cx, \quad Ax \geq b \in \mathbb{R}^m, \quad x \in \text{conv}(X) \quad (2)$$

or its Dantzig-Wolfe formulation: the (linear) master problem

$$\min \sum_{i \in I} (cx^i) \lambda_i, \quad \sum_{i \in I} (Ax^i) \lambda_i \geq b, \quad \sum_{i \in I} \lambda_i = 1, \quad \lambda_i \geq 0, \quad i \in I. \quad (3)$$

Remark 1.1 *Needless to say, the points in $\text{conv}(X)$ have the form $\sum_i \lambda_i x^i$, with λ varying in the unit simplex.*

Our framework could handle the unbounded case as well, including generating rays c^j : $\text{conv}(X)$ would have the form $\sum_i \lambda_i x^i + \sum_j \mu_j c^j$, the μ_j being just nonnegative (not restricted to sum up to 1); X could also be a mixed-integer polyhedron; or also be the Cartesian product of several subsets, i.e. $X = \prod_\ell X^\ell$ (when the constraint matrix defining X has a block diagonal structure, say).

Some of these "extensions" shall be illustrated when we come to specific applications in Section 2. It is mainly for simplicity that here we restrict our presentation to a single-block bounded integer polyhedron X . \square

Our aim is to solve the Dantzig-Wolfe formulation by column generation. At the current step k of the process, a *restricted master problem* is solved, obtained by restricting I to some $I^k \subset I$ in (2) or (3); we will set correspondingly $X^k := \{x^i : i \in I^k\}$. This resolution provides essentially two outputs:

- a primal solution $\hat{x} = \sum \hat{\lambda}_i x^i$, which is a candidate to solving (2) or (3),
- a dual solution $(u^k, r^k) \in \mathbb{R}_+^m \times \mathbb{R}$ associated with constraints $Ax \geq b$ and $\sum_i \lambda_i = 1$ respectively.

Then, the dual solution is used to price out columns in $I \setminus I^k$. For this, an optimisation subproblem or *oracle* is called upon to provide new columns of negative reduced cost, if any: for the given $u = u^k$, we minimize (possibly approximately) $(c - u^k A)x^i$ for i ranging over the whole of I , i.e. we consider the problem

$$\min_{x \in X} (c - uA)x. \quad (4)$$

at $u = u^k$, where X can equally be replaced by $\text{conv}(X)$.

From now on we will assume for simplicity that the oracle is exact and that $I^k = \{1, 2, \dots, k\}$: for a given u , the oracle solves (4) exactly, and provides just one optimal column, call it $x(u)$.

1.1 Standard column generation

The traditional – and most natural – restricted master problem is just (2) or (3) with I replaced by I^k : we solve

$$\min cx, \quad Ax \geq b \in \mathbb{R}^m, \quad x \in \text{conv}(X^k), \quad (5)$$

which is the linear program

$$\min \sum_{i=1}^k \lambda_i cx^i, \quad \sum_{i=1}^k \lambda_i Ax^i \geq b, \quad \sum_{i=1}^k \lambda_i = 1, \quad \lambda_i \geq 0, \quad i = 1, \dots, k. \quad (6)$$

Note that feasibility must be enforced: the initial (5) or (6) must contain a number of columns (possibly artificial) having in their convex hull some x satisfying $Ax \geq b$.

Suppose (5) or (6) has indeed a solution \hat{x} or $\hat{\lambda}$ with $\hat{x} = \sum \hat{\lambda}_i x^i$, and multipliers (u^k, r^k) . To check whether \hat{x} solves (2), we implicitly do full pricing by solving (4) to obtain $x(u^k)$. Two cases can arise:

- (i) Either $x^{k+1} := x(u^k)$ has negative reduced cost: $(c - u^k A)x^{k+1} + r^k < 0$; this new column is appended to the “bundle” x^1, \dots, x^k and the process is repeated, (note: with our definition of (u^k, r^k) , a negative reduced cost implies that x^{k+1} does not lie in $\text{conv}(X^k)$).
- (ii) Or $x(u^k)$ has zero reduced cost: then all columns in I have nonnegative reduced cost and (2) or (3) is solved.

Remark 1.2 *Some tolerance can be inserted, to stop the algorithm when the minimal reduced cost is “not too negative”, i.e. $(c - u^k A)x^{k+1} + r^k \geq -\varepsilon$.* \square

Insofar as we aim at solving (2), \hat{x} plays its role, of course. However, it is important to understand that the key role is actually played by u^k which, via the column x^{k+1} , really drives the algorithm toward convergence. Good u^k 's are those that are close to optimal multipliers for (2) or (3). It is therefore of interest to analyze (5) or (6) in the dual space, the LP dual of (6) being the “dual restricted master”

$$\max bu - r, \quad uAx^i - r \leq cx^i, \quad i = 1, \dots, k, \quad (u, r) \in \mathbb{R}_+^m \times \mathbb{R}. \quad (7)$$

1.2 Column generation seen in the dual space

For subsequent use, it is convenient to free oneself from the LP formalism and to use general duality (see for example [10], [16, Chap. XII], [29, §1], [28, §1.2]).

To (2) we associate the Lagrange function

$$\text{conv}(X) \times \mathbb{R}^m \ni (x, u) \mapsto L(x, u) := cx + u(b - Ax) = (c - uA)x + bu \quad (8)$$

and the so-called *dual function*

$$\begin{aligned} \mathbb{R}^m \ni u \mapsto \theta(u) &:= \min_{x \in X} L(x, u) = \min_{x \in \text{conv}(X)} L(x, u) \\ &= (c - uA)x(u) + bu; \end{aligned} \quad (9)$$

in the last expression, recall that $x(u)$ denotes the answer of the oracle, i.e. an argmin in (4). The dual problem associated with (2) is then

$$\max \{ \theta(u) : u \in \mathbb{R}_+^m \}, \quad (10)$$

which is the linear program

$$\max_{(u, s) \in \mathbb{R}_+^m \times \mathbb{R}} \{ s : s \leq bu + (c - uA)x^i, \quad i \in I \}. \quad (11)$$

Performing the same operations on the restricted master, we introduce likewise the *restricted dual function*

$$\mathbb{R}^m \ni u \mapsto \theta^k(u) := \min_{x \in \text{conv}(X^k)} L(x, u) = bu + \min_{i=1}^k (c - uA)x^i, \quad (12)$$

which overestimates θ : $\theta^k(u) \geq \theta(u)$ for all u . The dual of (5) is then

$$\max \{ \theta^k(u) : u \in \mathbb{R}_+^m \}, \quad (13)$$

i.e.,

$$\max_{(u, s) \in \mathbb{R}_+^m \times \mathbb{R}} \{ s : s \leq bu + (c - uA)x^i, \quad i = 1, \dots, k \}. \quad (14)$$

Setting $s = bu - r$, we recognize (7) in (14): s [resp. r] stands for $\theta^k(u)$ [resp. $bu - \theta^k(u)$]; r can be understood as a fixed cost in the primal oracle (4). Observe here and now that u solves (13) when $(u, \theta^k(u))$ solves (14), or equivalently when $(u, bu - \theta^k(u))$ solves (7).

Column generation for (6) is row generation for its dual (7). Thus, the algorithm of §1.1 seen in the dual is a cutting-plane procedure to maximize θ . At each iteration, we maximize θ^k instead, as illustrated on the right part of Fig. 1. Then we can write

$$\theta(u) \leq \theta^k(u) \leq \theta^k(u^k) \quad \text{for all } u \in \mathbb{R}_+^m. \quad (15)$$

To check whether the iterate u^k thus obtained does maximize θ , we compute the true value $\theta(u^k)$: we call the oracle (4) or (9) to obtain $x(u^k)$. Said otherwise, we check whether u^k is feasible for the unrestricted dual master (11), calling on (4) or (9) to identify the most violated inequality. Two cases can arise:




Figure 1: Standard column generation = Kelley-Cheney-Goldstein

- (i) Either $\theta(u^k) < \theta^k(u^k)$; then, the new cut $x^{k+1} := x(u^k)$ is appended to the bundle and the process is repeated.
- (ii) Or $\theta(u^k) = \theta^k(u^k)$ and (15) shows that u^k maximizes θ .

The above operations are clearly the same as (i), (ii) in §1.1. Note that equality between primal and dual optimal values of (6) and (7) gives $\sum_i \hat{\lambda}_i c x^i = b u^k - r^k$, or $c \hat{x} = \theta^k(u^k)$.

Remark 1.3 *The algorithm can be stopped when the largest cut violation is “sufficiently small”, i.e. when*

$$(u^k A - c)x^{k+1} - r^k = \theta^k(u^k) - \theta(u^k) = c \hat{x} - \theta(u^k) \leq \varepsilon,$$

indicating ε -optimality of both \hat{x} and u^k . Observe that the most negative reduced cost of Remark 1.2, the maximum violation and the duality gap $c \hat{x} - \theta(u^k)$ are all the same number. Thus, the tolerance ε is measured in the same units as the objective function of (2) and can be conveniently controlled.

However, because the sequence $\theta(u^k)$ is not monotone, it is more efficient to use the best dual iterate

$$\hat{u} := \text{Argmax} \{ \theta(u^i) : i = 1, \dots, k \}. \quad (16)$$

If

$$c \hat{x} - \theta(\hat{u}) \leq \varepsilon, \quad (17)$$

then \hat{x} and \hat{u} are ε -optimal for the primal and dual problems respectively. \square

Primal-dual relations coming from standard LP theory can be derived in our convex-analysis language. The set of optimal solutions in (12) (for given u), namely

$$\hat{X}^k(u) := \{ x \in \text{conv}(X^k) : L(x, u) = \theta^k(u) \} \quad (18)$$

is important for this. If u solves (13), $\hat{X}^k(u)$ gathers all columns associated to active constraints in (7) or (14), or equivalently with zero reduced cost in (6).

Theorem 1.4

- (i) *An optimal solution u^k of the restricted master (13) is characterized by the existence of some $\hat{x} \in \hat{X}^k(u^k)$ satisfying complementarity slackness:*

$$A \hat{x} - b \geq 0, \quad u^k \geq 0, \quad (A \hat{x} - b) u^k = 0$$

(abbreviated as $0 \leq A \hat{x} - b \perp u^k \geq 0$).

- (ii) *The \hat{x} 's described in (i) coincide with the optimal solutions of (5).*
- (iii) *With \hat{u} defined in (16), these \hat{x} 's are optimal in (2) if $c \hat{x} = \theta(\hat{u})$.*

Proof. Elementary convex analysis says that θ is concave and its subdifferential (the set of $g \in \mathbb{R}^m$ such that $\theta^k(v) \leq \theta^k(u) + g(v - u)$ for all $v \in \mathbb{R}^m$) is

$$\partial \theta(u) = \cup \{ g = b - Ax : x \in \hat{X}^k(u) \};$$

see for example [17, §D4.3-4] or [16, §VI.4.3-4]. Then u^k maximizes θ^k if and only if some subgradient lies in the normal cone to \mathbb{R}_+^m at u^k ; this is (i).

For (ii), invoke for example [42, Thm. 28.1] or [16, Thm. VII.4.4.3]: the pairs (\hat{x}, u^k) solving the pair of dual problems (5), (13) are the saddle-points of the Lagrangian (8), i.e. those feasible pairs satisfying

$$\forall (x, u) \in \text{conv}(X^k) \times \mathbb{R}_+^m, L(x, u^k) \geq L(\hat{x}, u^k) \geq L(\hat{x}, u).$$

It is not difficult to check that this repeats (i).

As for (iii), set $\varepsilon = 0$ in (17). □

Let us conclude:

- Standard column generation consists in merely restricting X to X^k in (2), leaving everything else unchanged. In the dual space, θ is correspondingly “impoverished” to θ^k . In nonlinear optimization, this is known as the cutting-plane method of Kelley [18] or Cheney-Goldstein [5].
- It needs a number of columns to start.
- An example due to A.S. Nemirovskii ([37, §4.3.6], reproduced in [16, §XV.1.1]) shows that as many as $(1/\varepsilon)^{m/2}$ calls to the oracle may be necessary to reach accuracy ε . Thus, standard column generation can be desperately slow.
- On the other hand, it directly provides a feasible primal point \hat{x} which satisfies complementarity slackness with u^k and is a distinguished candidate to solving (2).
- As a result, it can be safely stopped as soon as the window $c\hat{x} - \theta(u^k)$ or $c\hat{x} - \theta(\hat{u})$ is small.

To finish this section, note that the method of subgradient [47, 7, 40] is also a valid candidate to maximize θ , as well as the ellipsoid method [44, 36], which reaches accuracy ε after $(\log 1/\varepsilon)^m$ calls. Both methods behave poorly in practice, though. We also mention a little-known fact: the method of subgradients does provide a substitute for \hat{x} , as well as a stopping criterion resembling (17); see [45, 1, 25].

1.3 Stabilized column generation

Seen in the dual space, where the problem is to maximize θ , the rationale for (7) is the hope that the restriction to a subset of k columns gives a good approximation (12) of the true dual function (9): $\theta^k \simeq \theta$, so that maximizing θ^k should do good in terms of maximizing θ . Figure 2 reproduces the right part of Fig. 1, in which the horizontal line is the level $\theta(\hat{u})$ of the best answer of the oracle (16). Because $\theta^k \geq \theta$, the optimal $(u^*, \theta(u^*))$ clearly lies somewhere in the “safeguard” polyhedron P of Fig. 2. Standard column generation chooses the highest point in P .




Figure 2: The safeguard polyhedron

Nemirovskii’s example warns us that choosing this highest point may be far too optimistic and we may well have $\theta(u^k) \ll \theta^k(u^k)$. It is advisable to adopt more conservative strategies, known as *stabilized column generation*. Instead of θ^k , one maximizes some other function, providing some less high but more central point in P . In most of these methods, the dual restricted master (13) is replaced by

$$\max_{u \in \mathbb{R}_+^m} \tilde{\theta}^k(u) \quad \text{where} \quad \tilde{\theta}^k(u) := \theta^k(u) - S(u - \hat{u}). \quad (19)$$

Here the stability center \hat{u} is defined in (16); the stabilizing function S should have 0 as a minimum point, in order to pull u^k toward \hat{u} . Up to some minor variations, these methods work as follows.

Algorithm 1.5 (Schematic stabilized column generation)

STEP 0. Choose an initial stability center \hat{u} . Select an initial set of k columns.

STEP 1. Compute an optimal solution u^k of the stabilized dual restricted master problem (19).

STEP 2. Call the oracle (4) at u^k to obtain the new column x^{k+1} and the dual value $\theta(u^k) = L(x^{k+1}, u^k)$.

Perform the stopping test.

STEP 3. If $\theta(u^k) > \theta(\hat{u})$ set $\hat{u} = u^k$.

STEP 4. Increase k by 1 and go to Step 1. □

Step 3 is motivated by (16). The initial \hat{u} may be set to any heuristic estimate of a dual optimal solution to (3); as a default, \hat{u} may be initialized to 0. Most such algorithms may start at $k = 0$ with no initial column, u being initialized to \hat{u} in Step 1. Standard column generation uses $S \equiv 0$ and does need a nonempty set of initial columns.

Beware that, in the stabilized version, u^k no longer maximizes θ^k , so Thm. 1.4 does not apply. The role of that theorem was to allow the construction of the primal candidate \hat{x} from the optimality conditions in the dual restricted master (7), or as a solution of (6), viewed as the dual of (7). Stabilized duals of the form (19) will provide an alternative candidate – denoted by \tilde{x} hereafter. To obtain this candidate by solving an optimization problem, we need to define a dual of (19) resembling (5). This dualization is indeed possible via conjugate calculus (see [16, Chap. XI] or [17, Chap. E]). Specifically, introduce the *conjugate function*

$$S^*(g) := \max_{v \in \mathbb{R}^m} [gv - S(v)]; \quad (20)$$

as will be seen in Thm. 1.7 below, a convenient dual to (19), called the *Fenchel dual*, is then

$$\min cx + \hat{u}g + S^*(g), \quad Ax \geq b - g, \quad (x, g) \in \text{conv}(X^k) \times \mathbb{R}^m. \quad (21)$$

This is admittedly a rather abstract problem, it will be made more precise when S is specified.

Remark 1.6 A way to obtain (21) is to formulate (19) as

$$\max [\theta^k(u) - S(v)], \quad v = u - \hat{u}, \quad (u, v) \in \mathbb{R}_+^m \times \mathbb{R}^m$$

and to apply Lagrangian relaxation: associating a multiplier $g \in \mathbb{R}^m$ with the constraint $v = u - \hat{u}$, the Lagrangian $\theta^k(u) - S(v) + g(v - u + \hat{u})$ produces the dual

$$\min_{g \in \mathbb{R}^m} \left[\max_{u \geq 0} (\theta^k(u) - gu) + S^*(g) + \hat{u}g \right].$$

Tedious calculations show that this coincides with (21).

Also, we mention two situations where (21) simplifies:

– An equality-constrained master problem (5) produces $Ax = b - g$ in (21), which becomes

$$\min cx + \hat{u}(b - Ax) + S^*(b - Ax), \quad x \in \text{conv}(X^k).$$

The connection with (5) is more transparent; in particular, the above minimand discloses the “augmented Lagrangian” $L(x, \hat{u}) + S^*(b - Ax)$, with an augmenting function S^* . Remembering that S is minimal at 0, it can be shown that S^* is also minimal at 0, and therefore pulls $b - Ax$ toward 0.

– Standard column generation uses $S \equiv 0$, whose conjugate is clearly

$$S^*(g) = \begin{cases} 0 & \text{if } g = 0, \\ +\infty & \text{otherwise} \end{cases} =: i_{\{0\}}(g),$$

the so-called indicator function of $\{0\}$: the term $S^*(g)$ forces $g = 0$ in (21), which becomes exactly (5). □

We can now state an adapted version of Thm. 1.4:

Theorem 1.7 Assume that S is a convex function over \mathbb{R}^m . Consider an optimal solution u^k of (19) such that $S(v) < +\infty$ for all v close to $u^k - \hat{u}$. Then:

(i) There exist $\tilde{x} \in \hat{X}^k(u^k)$ of (18) and a subgradient $\tilde{g} \in \partial S(u^k - \hat{u})$ such that $0 \leq (A\tilde{x} - b + \tilde{g}) \perp u^k \geq 0$.

(ii) The (\tilde{x}, \tilde{g}) 's described in (i) coincide with the optimal solutions of (21).

(ii') If $\tilde{g} = 0$ in (i), then the corresponding \tilde{x} is optimal in (5).

(iii) If, in addition, $c\tilde{x} = \theta(\hat{u})$, then \tilde{x} is optimal in (2).

Proof. The local finiteness of S is just a technical assumption, to guarantee in (19) that $\partial\tilde{\theta}(u^k) = \partial\theta^k(u^k) - \partial S(u^k - \hat{u})$ (see [16, §XI.3.5]; other assumptions are possible, for example that $\tilde{\theta}$ is polyhedral). Then (i) goes as in Thm.1.4(i); (ii) is the so-called Fenchel duality theorem and the rest is clear. \square

Thus, \hat{x} of §1.1 is replaced by \tilde{x} , which can be obtained by proving optimality of u^k in (19), or by directly solving its associated “bidual” (21). Recall the definition (18) of $\tilde{X}^k(u^k)$: as before, $\tilde{x} = \sum_i \lambda_i x^i$ for a number of columns i such that $L(x^i, u^k) = \theta(u^k)$. What is missing is *feasibility* – and a fortiori complementarity: $A\tilde{x} \not\geq b$, unless $\tilde{g} \leq 0$. As a result, two things must be checked before a stabilized algorithm can be stopped: $c\tilde{x} - \theta(\hat{u})$ must be small as before, but the positive part of \tilde{g} must also be small (so that \tilde{x} is approximately feasible). See [9] for a convergence study of general stabilized methods.

Several choices have been proposed for S , which we briefly review now.

1.3.1 Boxstep

A first stabilizing device forces u in an ℓ_∞ -box around \hat{u} . The stabilizing problem is

$$\max \{ \theta^k(u) : u \geq 0, \|u - \hat{u}\|_\infty \leq \delta \}, \quad \text{i.e.} \quad \max_{u \in \mathbb{R}_+^m} \theta^k(u) - i_{B_\delta}(u - \hat{u}); \quad (22)$$

S is the indicator function of B_δ , see Fig.3.




Figure 3: The indicator function of B_δ

This is the boxstep method of [33]. Its original aim was not much to reduce the number of calls to the oracle, but rather to ease the resolution of (6), replacing it by

$$\max \theta(u), \quad \|u - \hat{u}\|_\infty \leq \delta, \quad u \in \mathbb{R}_+^m.$$

The stability center \hat{u} was chosen a priori; it was not moved as in (16) but only after the above problem was fully solved (unless $\|u^k - \hat{u}\|_\infty < \delta$, indicating overall optimality and termination of the column generation process).

The conjugate (20) is easy to compute when S is an indicator function: this gives $i_{B_\delta}^*(g) = \delta \|g\|_1$. According to Remark 1.6 and Thm.1.7, Boxstep therefore produces an \tilde{x} solving

$$\min cx + \hat{u}g + \delta \|g\|_1, \quad Ax \geq b - g, \quad (x, g) \in \text{conv}(X^k) \times \mathbb{R}^m. \quad (23)$$

Both problems (22) and (23) can be formulated as linear programs; this will be done in §1.3.3 below.

1.3.2 Polyhedral penalties

A series of stabilizing terms S have been subsequently proposed, in which the stabilized dual restricted master (19) is a more and more sophisticated LP. The corresponding stabilizing terms are depicted in Fig.4.




Figure 4: Stabilization by ℓ_1 penalty

In [19], S is V-shaped (left part of Fig.4): (19) has the form

$$\max_{u \in \mathbb{R}_+^m} \theta^k(u) - s \sum_{j=1}^m |u_j - \hat{u}_j|, \quad \text{i.e.} \quad \max_{u \in \mathbb{R}_+^m} \theta^k(u) - s \|u - \hat{u}\|_1, \quad (24)$$

with possibly $2m$ penalty coefficients s_j^+ and s_j^- instead of the single s . Interestingly enough, this stabilization is the dual counterpart of Boxstep: in fact, $\|\cdot\|_1^* = \|\cdot\|_\infty$ and the Fenchel bidual (21) takes the form

$$\min cx + \hat{u}g, \quad Ax \geq b - g, \quad \|g\|_\infty \leq s, \quad (x, g) \in \text{conv}(X^k) \times \mathbb{R}^m, \quad (25)$$

which can again be formulated as a linear program.

The penalty used in [35] has a pair of breakpoints (central part of Fig.4), again with possibly adaptable slopes s, s' and breakpoints δ' . The stabilizer of the right part of Fig.4 is proposed in [4], where the management of \hat{u} is inspired by boxstep, rather than (16).

A common feature of these stabilizations is the necessity to manage carefully their parameters s, δ , etc. For example:

- In the boxstep stabilization of Fig. 3, a large δ stabilizes nothing, while a small δ slows down the iterates unduly.
- The form S_1 of Fig. 4 requires $s \rightarrow 0$, to force a small \tilde{g} in Thm. 1.7 – see (25).

1.3.3 Polyhedral stabilization in the primal

The main merit of the stabilizations reviewed so far is to allow LP formulations, both in the primal and dual spaces – as was alluded to for (23) and (25). Let us make explicit the corresponding LP programs in the case $S(v) = s \sum_{j=1}^m \max\{0, |v_j| - \delta\}$.




Figure 5: An illustrating stabilization (and its conjugate)

Each one-dimensional component j of S (depicted on the left part of Fig. 5) can be represented by

$$S_j(v) = s \min\{\rho_j : \rho_j \geq v_j - \delta, \rho_j \geq -v_j - \delta\}.$$

As a result, the LP formulation of (19) is here

$$\begin{cases} \max bu - r - s \sum_{j=1}^m \rho_j, \\ uAx^i - r \leq cx^i & i = 1, \dots, k, \\ u_j - \rho_j \leq \hat{u}_j + \delta & j = 1, \dots, m, \\ -u_j - \rho_j \leq -\hat{u}_j + \delta & j = 1, \dots, m, \\ u, \rho \geq 0, \end{cases} \quad \begin{array}{l} \text{multipliers } \lambda_i \\ \text{multipliers } g_j^+ \\ \text{multipliers } g_j^- \end{array}$$

whose LP dual is (A_j being the j^{th} row of A)

$$\begin{cases} \min \sum_{i=1}^k (cx^i)\lambda_i + \sum_{j=1}^m (\delta + \hat{u}_j)g_j^+ + \sum_{j=1}^m (\delta - \hat{u}_j)g_j^- , \\ \sum_{i=1}^k \lambda_i = 1, \\ \sum_{i=1}^k (A_j x^i)\lambda_i + g_j^+ - g_j^- \geq b_j & j = 1, \dots, m, \\ g_j^+ + g_j^- \leq s & j = 1, \dots, m, \\ \lambda, g^+, g^- \geq 0. \end{cases} \quad (26)$$

Stabilization amounts to relaxing the constraints $Ax \geq b$ by introducing slack variables g_j^+ and g_j^- , whose sum is bounded by s , and with marginal costs $\delta + \hat{u}_j$ and $\delta - \hat{u}_j$ respectively.

Remark 1.8 *It can be seen that g_j^+ and g_j^- cannot be simultaneously positive: in fact, $u_j - \hat{u}_j$ is – either nonzero, then one of the constraints*

$$\rho_j \geq u_j - \hat{u}_j - \delta, \quad \rho_j \geq \hat{u}_j - u_j - \delta$$

is inactive,

– or 0, then one of them is redundant.

A consequence is that the constraints $g_j^+ + g_j^- \leq s$ can be replaced by $g_j^+ \leq s, g_j^- \leq s$. \square

LP formulations are possible because the penalty function S and its conjugate S^* can be written via linear relations. An additional illustration is then to derive the ‘‘Fenchel bidual’’ (21) on the above example (a rather clumsy approach, for want of a compact expression of S , though). Conjugating the S_j of Fig. 5 can be done by tedious calculations. A first key is to observe that

$$S_j(v_j) = s \max_{(\lambda, \mu) \in \Delta} [\lambda(v_j - \delta) - \mu(v_j + \delta)] = -s \min_{(\lambda, \mu) \in \Delta} [\lambda(\delta - v_j) + \mu(v_j + \delta)];$$

here $\Delta := \{(\lambda, \mu) : \lambda \geq 0, \mu \geq 0, \lambda + \mu \leq 1\}$. Then the min and max can be inverted in the expression

$$S_j^*(g_j) := \max_{v_j} \min_{(\lambda, \mu) \in \Delta} [(g_j + s(\mu - \lambda))v_j + s(\lambda + \mu)\delta],$$

which results in

$$S_j^*(g_j) = s\delta \min_{(\lambda, \mu) \in \Delta} \{\lambda + \mu : g_j + s(\lambda - \mu) = 0\}.$$

Working out the calculations gives the function depicted on the right part of Fig. 5. Then apply for example Prop.1.3.1(ix) of [16, Chap.X] or [17, Chap.E]: our Fenchel bidual (21) boils down to

$$\min cx + \hat{u}g + \delta\|g\|_1, \quad Ax \geq b - g, \quad \|g\|_\infty \leq s, \quad x \in \text{conv}(X^k).$$

Setting $s = +\infty$ [resp. $\delta = 0$] reproduces (23) [resp. (25)].

Now set $g := g^+ - g^- \in \mathbb{R}^m$ in (26) and use Remark 1.8: we clearly have $g_j^+ + g_j^- = |g_j|$. Thus,

$$\|g\|_1 = \sum_{j=1}^m (g_j^+ + g_j^-), \quad \|g\|_\infty = \max_{j=1}^m (g_j^+ + g_j^-),$$

and we find back (26).

1.3.4 Euclidean penalty: bundle methods

Contemplating the sequence of stabilizers of Fig. 4 suggests that they tend to mimic a quadratic function – which takes δ and s ‘‘infinitely small’’, and ‘‘infinitely many’’ breakpoints. In fact, Kelley’s method is not the only one to need stabilization: other methods in nonlinear programming also need it, and there are good reasons to use quadratic stabilizers; see some explanations in [16, §II.2.2(c)].

Bundle methods go back to [26]. In their latest form, initiated in [27, 21] and fully described in [16, Chap. XV], they use a Euclidean norm for S in (19): u^k solves

$$\max_{u \in \mathbb{R}^m} \theta^k(u) - \frac{1}{2t} \|u - \hat{u}\|^2. \quad (27)$$

Clearly, $S^*(g) = \frac{t}{2} \|g\|^2$, so the Fenchel bidual (21) writes

$$\min cx + \hat{u}g + \frac{t}{2} \|g\|^2, \quad Ax \geq b - g, \quad (x, g) \in \text{conv}(X^k) \times \mathbb{R}^m. \quad (28)$$

Solving this problem with respect to g for fixed x gives the explicit value $g = \max\{b - Ax, -\hat{u}/t\}$; plugging it back into the objective function discloses the augmented Lagrangian of [43]. A rather compact writing of the resulting problem is

$$\min_{x \in \text{conv}(X^k)} cx + \frac{1}{2t} \|(t(b - Ax) + \hat{u})_+\|^2 - \frac{1}{2t} \|\hat{u}\|^2.$$

Remark 1.9 Assume an equality-constrained master problem (5), in which case u is unconstrained in (27). As mentioned in Remark 1.6, the Fenchel bidual simplifies to

$$\min L(x, \hat{u}) + \frac{t}{2} \|b - Ax\|^2, \quad x \in \text{conv}(X^k). \quad (29)$$

The augmented Lagrangian is now more transparent: instead of imposing the constraint $Ax = b$ as in (5), we

- dualize it via $\hat{u}(b - Ax)$ (observe the careful choice of the dual variable, which is set to the stability center \hat{u}),
- and penalize it via $\frac{t}{2}\|b - Ax\|^2$, which limits constraint violations by the candidate \tilde{x} . \square

The relevant primal-dual relations now rely on a crucial output of the quadratic restricted master (27):

$$\tilde{g} := \frac{u^k - \hat{u}}{t}. \quad (30)$$

Theorem 1.10 Use the notation (30).

- (i) The unique optimal solution u^k of (27) is characterized by the existence of some $\tilde{x} \in \hat{X}^k(u^k)$ such that $0 \leq (A\tilde{x} - b + \tilde{g}) \perp u^k \geq 0$.
- (ii) Together with \tilde{g} , the \tilde{x} 's described in (i) coincide with the optimal solutions of (28).
- (ii') If $u^k = \hat{u}$ then these \tilde{x} 's coincide with the optimal solutions of (5).
- (iii) In case (ii'), they are also optimal in (2).

Proof. That (27) has a unique optimal solution is clear. Then apply Thm. 1.7. Here (21) is (28) and $\partial S(u - \hat{u})$ is clearly the singleton $(u - \hat{u})/t$, providing the explicit value for $\tilde{g} = (u^k - \hat{u})/t$.

In case (ii'), we have $0 \in \partial\theta^k(u^k) - 0$, hence u^k maximizes θ^k . Besides, $\tilde{g} = 0$ in (i) shows that \tilde{x} is feasible in (5) and (2). Finally, observe in Alg. 1.5 that $\theta^k(\hat{u}) = \theta(\hat{u})$. Thus, \hat{u} maximizes θ and (iii) is proved. \square

The proof of (iii) deserves a comment. First observe that a key is the existence of the gradient $\tilde{g} = \nabla S(u^k - \hat{u})$. Indeed (iii) holds in Thm. 1.7 as well, whenever S in (19) is a differentiable function. Also, (iii) seems paradoxical: it implies that checking $c\tilde{x} = \theta(\hat{u})$ is unnecessary to stop the process. The trickery here is that the oracle has actually been already called at \hat{u} , during some previous iteration. This guarantees the last equality in the chain

$$\forall u \in \mathbb{R}^m, \theta(u) \leq \theta^k(u) \leq \theta^k(u^k) = \theta^k(\hat{u}) = \theta(\hat{u}).$$

The crucial \tilde{g} of (30) measures the “default of complementarity” of \tilde{x} , or the “default of optimality” of \hat{u} ; and its convergence to 0 conditions the convergence of a bundle method. This latter property does not depend much on t : it holds for example with $t \equiv 1$. However practical efficiency does require a $t = t^k$ adapted at each iteration.

1.3.5 ACCPM

The Analytic-Center Cutting-Plane Method of [11] replaces θ^k by a barrier function as follows. With respect to the variable $(u, s) \in \mathbb{R}_+^m \times \mathbb{R}$, the safeguard polyhedron P of Fig. 2 is defined by the k constraints of the dual restricted master (14), and a level constraint, namely

$$\ell_i(u, s) := L(x^i, u) - s \geq 0, \quad i = 1, \dots, k, \quad \text{and} \quad \ell_0(u, s) := s - \theta(\hat{u}) \geq 0.$$

The analytic center of P is then the unique point maximizing the barrier function $B(u, s) := \sum_{i=0}^k \log \ell_i(u, s)$. Note that replacing θ^k by B is a stabilization in itself: the analytic center certainly satisfies the constraints strictly (see Fig. 2 again). The constraints ℓ_i are actually suitably scaled, and a complexity analysis [12] shows that a few Newton iterations maximizing B suffice to guarantee good overall convergence. In a variant of the method [38, 39], an extra quadratic stabilizer $\frac{1}{2t}\|u - \hat{u}\|^2$ is subtracted from B and convergence is improved; see also [13].

2 Numerical comparisons

We have compared numerically the bundle stabilization of §1.3.4 with the standard Kelley algorithm of §1.1. Ideally, a comparison with the most advanced linear programming stabilization of Kelley would be interesting. However, the necessity of application-specific tuning of various parameters makes it impractical. The applications used for these comparative tests are the cutting-stock, the multi-item lot-sizing, the traveling-salesman, the capacitated vehicle-routing, and the vertex-coloring problems. They illustrate the diversity of models that can arise when X has a more complex structure, as alluded in Remark 1.1. Each application is introduced by presenting the specific form taken by the original integer program, the associated Lagrangian dual and its Dantzig-Wolfe formulation. First we give some implementation details.

2.1 Implementations

The algorithms described in §1.1 and §1.3.4 are implemented using the environment of *BaPCod*, a generic Branch-And-Price Code [49]. Xpress^{MP} [6] is used for solving master linear programs while customized solvers are used for the oracles. We emphasize that, in all of our tests below, the oracle minimizes the Lagrangian exactly (returns the most negative reduced cost), even though this may not be the usual practice in column generation.

Kelley

Two versions are considered: **KBASIC** and **KRICH**, which differ by their initial columns. These “columns”, only characterized by their cost and constraint value, are purely artificial: they do not correspond to any element of X , or even of $\text{conv}(X)$. As a result, they must be eliminated from the solution \hat{x} of the master before a reliable feasible cost $c\hat{x}$ can be used for the stopping test.

- In the basic version **KBASIC**, the algorithm of §1.1 is initialized with a single artificial column – denote it by x^1 – with constraint value $Ax^1 \geq b$, and whose cost cx^1 is set to an estimate of the optimal value of (1).
- **KRICH** is initialized with m artificial columns. Their constraint values form the canonical basis of \mathbb{R}^m : A_j being the j^{th} row of A , $A_j x^i = 1$ if $i = j$, 0 otherwise; and $cx^i = \alpha \hat{u}_i$ where \hat{u} is an application-specific heuristic estimate of the optimal dual solution of (2) and α is a factor, set to 1.2 unless otherwise specified¹. For both versions, the algorithm is stopped when (17) holds. However, if artificial columns are in the primal solution $\hat{\lambda}$ at that stage, their cost is increased by the factor α and the column generation procedure is continued.

Remark 2.1 Let $j = 1, \dots, m$ index the artificial columns in **KRICH** and call g_j their multipliers: the k^{th} primal master is

$$\begin{cases} \min \sum_{i=1}^k (cx^i) \lambda_i + \alpha \sum_{j=1}^m \hat{u}_j g_j, \\ \sum_{i=1}^k \lambda_i = 1, \\ \sum_{i=1}^k A_j x^i \lambda_i + g_j \geq b_j & j = 1, \dots, m, \\ \lambda, g \geq 0. \end{cases}$$

This resembles (26), with g^- fixed to 0, $\delta_j = (\alpha - 1)\hat{u}_j$; also s is set to $+\infty$. The dual of the above program is

$$\begin{cases} \max bu - r, \\ uAx^i - r \leq cx^i & i = 1, \dots, k, \\ u_j \leq \alpha \hat{u}_j & j = 1, \dots, m, \\ u \geq 0. \end{cases}$$

Thus, **KRICH** resembles a sort of “static half-boxstep”, in which \hat{u} remains fixed and u is bounded from above only (0 being a natural lower bound). The size of the box is increased when the artificial cost is increased (to force the artificial variable out of the optimal LP solution). \square

Bundle

Our implementation follows rather accurately [32], in particular for the management of the stabilizing parameter $t = t^k$. The restricted master problem (27) is solved by the QP solver of K.C. Kiwiel [20, 23]. Unless otherwise specified, we start with an empty initial set of columns. Most of the time, two versions are tested: **BUNDLE**, where the initial iterate \hat{u} is problem-dependent, and **BUNDLE0** initialized at $\hat{u} = 0$. The algorithm needs an initial value for t , and this is obtained from an estimate of the optimal value of (2). The algorithm stops when (17) holds, together with $\|\tilde{g}\| \leq \varepsilon\sqrt{m}$ (meaning that \hat{x} satisfies the dualized constraints within ε on the average). The same value is used for ε throughout, for bundle and for Kelley.

Feasible primal points

Once the artificial columns are properly eliminated, Kelley’s method provides a feasible candidate \hat{x} at each iteration; the bulk of the work is then to improve this feasible solution. By contrast, primal feasibility is only asymptotic for a bundle method, which stops as soon as a feasible enough primal point is obtained – in a way, this is predicted by Theorem 1.10(iii). Reaching primal feasibility is thus the whole business of a bundle method. We have therefore considered **BUNDLE+K**, a variant aimed at combining the two approaches: bundle generates

¹The artificial x^i ’s are actually *rays*: it is only large enough multiples of them that contribute feasibility.

the sequence of dual iterates u^k , while Kelley is used to obtain a primal solution \hat{x}' based on which we perform an additional test at each iteration, in the hope to stop the algorithm earlier. Thus, BUNDLE+K is as follows (compare with Algorithm 1.5).

Algorithm 2.2 (Combination of bundle and Kelley)

STEP 0. Choose an initial stability center \hat{u} . Select an initial set of k columns.

STEP 1. Compute the optimal solution u^k of the quadratic restricted master problem (27).

STEP 1'. Solve the linear restricted master problem (6) to obtain \hat{x}' . Stop if $c\hat{x}' - \theta(\hat{u}) \leq \varepsilon$.

STEP 2. Call the oracle (4) at u^k to obtain the new column x^{k+1} and the dual value $\theta(u^k) = L(x^{k+1}, u^k)$.

Perform the stopping test.

STEP 3. If $\theta(u^k) > \theta(\hat{u})$ set $\hat{u} = u^k$.

STEP 4. Increase k by 1 and go to Step 1. □

Our numerical experiments have been performed on a PC pentium 3, 1 Ghz. Most of the tables below record the number of iterations and CPU times.

- One iteration means one execution of the oracle and of the master. Besides the number of iterations, the tables also record the total number of columns generated (including the artificial one for KBASIC, but not for KRICH and including the last subproblem solution with zero reduced cost when stopping does not arise before that). Note that, for a stabilized algorithm such as bundle, the oracle may answer the same column to several calls, even though these calls are made with different u 's.
- The total CPU time is spent respectively in the oracle and in the master, plus an overhead which may not be negligible. Note that when the oracle is a NP-Hard problem (as is the case for all applications but the multi-item lotsizing and the TSP), its computing time may vary in fairly large proportions.

2.2 The cutting stock problem (CSP)

The problem is to minimize the number of stock pieces of width W , used to meet demands d_1, \dots, d_m , for items $j = 1, \dots, m$, to be cut at their width w_1, \dots, w_m . We assume that every w_j is smaller than W and that there are enough stock pieces, say N , available for a feasible cutting: for example $N = \sum_j d_j$ is certainly enough.

Let us put the resulting problem in the framework of the present paper. Call $y^\ell \in \{0, 1\}$ an indicator of whether stock piece ℓ is used, and $z_j^\ell \in \mathbb{N}$ the quantity of item j cut in stock piece ℓ . Then, a possible formulation for (1) is

$$\begin{cases} \min \sum_{\ell=1}^N y^\ell, \\ \sum_{\ell=1}^N z_j^\ell \geq d_j & \text{for } j = 1, \dots, m, \\ \sum_{j=1}^m z_j^\ell w_j \leq W y^\ell & \text{for } \ell = 1, \dots, N, \\ y^\ell \in \{0, 1\}, z_j^\ell \in \mathbb{N} & \text{for } j = 1, \dots, m, \ell = 1, \dots, N. \end{cases} \quad (31)$$

Introducing the variable vector $x = (y, z)$, call $Ax \geq b$ the demand-covering constraints $\sum_\ell z^\ell \geq d$ and put the other constraints in X ; the Lagrangian (8) is then

$$L(y, z, u) = \sum_{\ell=1}^N (y^\ell - u z^\ell) + ud.$$

It is decomposable with respect to ℓ and its minimization produces the dual function (9) which takes the form

$$\theta(u) = du + \sum_{\ell=1}^N \min \{ y^\ell - u z^\ell : w z^\ell \leq W y^\ell, y^\ell \in \{0, 1\}, z^\ell \in \mathbb{N}^m \}.$$

This is the juxtaposition of N identical optimization problems, which can be solved by inspection on y^ℓ : one solves the knapsack problem²

$$\max uz, \quad wz \leq W, \quad z \in \mathbb{N}^m \quad (32)$$

²The knapsack solver proceeds by transforming the problem into a multiple-class binary knapsack problem and solves it using an extension of the standard branch-and-bound algorithm of Horowith and Sahni, as presented in [48].

to obtain an optimal solution $z(u)$ – we drop the useless index ℓ . The minimum value in the above curly bracket is then $\min\{0, 1 - uz(u)\}$, so that the dual function has the value

$$\theta(u) = du + \begin{cases} 0 & \text{if } uz(u) \leq 1, \\ N(1 - uz(u)) & \text{otherwise.} \end{cases} \quad (33)$$

An alternative presentation follows the Dantzig-Wolfe model (3). Let I enumerate solutions $x = (y, z)$, where z is feasible for (32) and y is the associated indicator variable. The problem can be formulated as

$$\min \sum_{i \in I} y^i \lambda_i, \quad \sum_{i \in I} z^i \lambda_i \geq d \in \mathbb{R}^m, \quad \sum_{i \in I} \lambda_i \leq N, \quad \lambda_i \geq 0, \quad i \in I.$$

Here $\sum_i \lambda_i \leq N$ is inherited from the original convexity constraints $\sum_i \lambda_i^\ell = 1$ associated with each stock piece: as stock pieces are identical, convexity constraints can be aggregated; the resulting constraint can be relaxed to a \leq -constraint since $(y, z) = 0$ is a solution of I . They express that the total number of available stock pieces is limited. Our N being an overestimate of the actual number of stock pieces used, this constraint is not binding and can be dropped: the formulation becomes

$$\min \sum_{i \in I} y^i \lambda_i, \quad \sum_{i \in I} z^i \lambda_i \geq d \in \mathbb{R}^m, \quad \lambda_i \geq 0, \quad i \in I, \quad (34)$$

whose dual is

$$\max du, \quad uz^i \leq 1, \quad i \in I, \quad u \in \mathbb{R}_+^m. \quad (35)$$

Remark 2.3 *It is interesting to observe that the formulation (35) of the dual is actually a nonlinearly constrained problem, namely*

$$\max du, \quad h(u) \leq 1, \quad u \geq 0,$$

where

$$h(u) := \max_{wz \leq W} uz = uz(u).$$

In other words, the oracle delivers the value of a dual constraint function h rather than a dual objective function θ . From this point of view, maximizing (33) appears as a so-called exact penalty approach to solve (35): the term $N(1 - uz(u))$ imposes a price to pay for infeasible u 's. \square

For this application, we compare in Tables 1 and 2 the two versions of Kelley, the standard bundle method, and the version BUNDLE+K of Algorithm 2.2. For KBASIC, the cost of the unique artificial column is set to $\lceil \frac{\sum_i w_i d_i}{W} \rceil$, a lower bound on the required number stock pieces. For KRICH, the artificial column costs are set to $\alpha \hat{u}$ with $\alpha = 1.2$ and $\hat{u} := w/W$. This value of \hat{u} solves the dual of the LP relaxation of (31); it turns out to be a very good estimation of the optimal solution to the Lagrangian dual. As for bundle, we tested two initializations:

- BUNDLE $\frac{1}{m}$: $\hat{u} = 1/m$ to test a poor initialization;
- BUNDLE: $\hat{u} = w/W$, an accurate initialization.

For BUNDLE+K, $\hat{u} = w/W$, and the master used for Kelley includes the same artificial column as in KBASIC. We also report on tests where the algorithm is initialized with the columns of an initial incumbent primal solution, obtained via a first-fit decreasing heuristic. These versions are denoted by Kelley-Inc, BUNDLE-Inc, and BUNDLE+K-Inc. Kelley-Inc is initialized with the columns of the heuristic solution only (no artificial columns are used). The counter “cols” in the tables includes these heuristic columns.

Also, the stopping test of Remark 1.3 is strengthened: knowing that the optimal value to (31) is integer, the dual bound is rounded-up to the next integer; earlier stop may occur when a primal bound is available (note: such is the case with Kelley but not with bundle).

Table 1 concerns the results on 9 industrial instances, whose number m of items ranges from 7 to 33; N is specified in the data. Table 2 gives the same statistics on 20 randomly generated instances, having 50 items each; the average demand is 100 and the widths are drawn uniformly in $[500, 5000]$, while the wideroll width is 10000 and $N = 2500$. The fastest methods seem to be KBASIC and Kelley-Inc, while the fewer iterations are with Kelley-Inc and BUNDLE+K-Inc. An important observation is that the number of calls to the oracle and the total time spent in the oracle are somehow anti-correlated: more time is spent in the oracle when it is called by a stabilized method. In fact, the more expensive oracles are those called by the end of the algorithms, when u and w tend to be collinear vectors: this behaviour is illustrated by Fig. 6, obtained with one of the random

Master	Counters		Timers (in seconds)		
	iter	cols	oracle	master	total
KBASIC	42	43	.38	.03	.52
KRICH	29	29	2.82	.01	2.91
BUNDLE $\frac{1}{m}$	90	52	6.86	.07	7.04
BUNDLE	43	33	42.57	.03	42.68
BUNDLE+K	38	32	38.20	.05	38.37
Kelley-Inc	24	49	.12	.03	.25
BUNDLE-Inc	41	56	40.47	.40	40.63
BUNDLE+K-Inc	25	49	6.33	.13	6.55

Table 1: CSP: averages on 9 industrial instances with 7 to 33 items

Master	Counters		Timers (in seconds)		
	iter	cols	oracle	master	total
KBASIC	235	236	3.45	.50	4.51
KRICH	142	142	3.15	.21	3.72
BUNDLE $\frac{1}{m}$	246	185	12.30	.54	13.37
BUNDLE	163	111	26.51	.27	27.17
BUNDLE+K	155	109	25.16	.39	26.08
Kelley-Inc	136	211	3.36	.20	4.61
BUNDLE-Inc	164	186	26.94	.40	28.70
BUNDLE+K-Inc	131	170	20.18	.46	22.00

Table 2: CSP: averages on 20 random instances with 50 items




Figure 6: CSP: oracle computing time along iterations

instances. A final observation is that the times spent in the master are (small and) comparable for all variants, whether this master is the LP (6) or the QP (28).

In another series of runs, we consider a variant of CSP: the number of items j to cut must take a value within a prescribed interval $[\underline{d}_j, \bar{d}_j]$. The objective is then to minimize the waste (this would be equivalent to minimizing the number of stock pieces if the demands were fixed). The resulting master takes the form

$$\min \sum_{i \in I} (W - wz^i) \lambda_i, \quad \underline{d} \leq \sum_{i \in I} z_j \lambda_i \leq \bar{d}, \quad \lambda_i \geq 0, \quad i \in I.$$

In this variant, convergence is no longer truncated by rounding dual bounds (the objective value is no longer an integer). The test problems are the same, but with a 5% tolerance on production: $\underline{d}_j = \lfloor 0.95 d_j \rfloor$ and $\bar{d}_j = \lceil 1.05 d_j \rceil$.

A possible heuristic to define \hat{u} is as follows. Let R be an upper bound on the optimal waste (obtained by solving the first variant). Then:

- set $u_j = \frac{Rw_j}{\sum_k w_k d_k}$ for the covering constraints $\sum_i z^i \lambda_i \geq \underline{d}$;
- set $u_j = \frac{R}{w_j \sum_k \frac{d_k}{w_k}}$ for the packing constraints $\sum_i z^i \lambda_i \leq \bar{d}$.

Intuitively, the portion of the waste induced by a covering constraint is proportional to the width of the item, while the waste induced by the production upper bounds is inversely proportional to the usefulness of the item in filling the holes (the more useful items being the small ones).

The results for this CSP variant are given in Tables 3 and 4. In these particular experiments, the QP (28) was cold-started at each iteration, even though its $k + 1^{\text{st}}$ execution could take advantage of the k^{th} one; a comparison of the master timers with the previous tables illustrates the benefit of warm starts. For KRICH, we define artificial column cost as $\alpha \hat{u}$ with $\alpha = 1.1$. The initialization BUNDLEW&0 refers to using \hat{u} as defined above for covering constraints but using 0 for packing constraints. For BUNDLE and BUNDLE+K, we use \hat{u} as defined above. To obtain an initial incumbent, the first-fit decreasing heuristic is used to cover demands \underline{d}_j ; then production surplus are used to fill holes in existing cutting patterns. Now the fastest method is KBASIC. Again, subproblems take longer to solve with stabilized schemes). Nevertheless, the fewer number of iterations is achieved by KRICH. Whatever the initialization, bundle takes more iterates than Kelley.

Master	Counters		Timers (in seconds)		
	iter	cols	oracle	master	total
KBASIC	43	43	2.39	.03	2.53
KRICH	31	26	8.42	.03	8.53
BUNDLE	60	35	15.35	.25	15.71
BUNDLEW&0	48	34	54.19	.23	54.51
BUNDLE+K	59	36	15.22	.30	15.66
Kelley-Inc	44	69	4.78	.05	4.97
bundle-Inc	55	59	19.28	.30	19.74
BUNDLE+K-Inc	54	59	19.35	.37	19.87

Table 3: CSP, minimization of the waste: averages on 9 industrial instances

Master	Counters		Timers (in seconds)		
	iter	cols	oracle	master	total
KBASIC	144	143	8.25	.29	9.03
KRICH	142	120	32.80	.29	33.56
BUNDLE	200	86	58.92	5.78	65.15
BUNDLEW&0	203	177	56.83	7.12	65.52
BUNDLE+K	198	87	58.20	5.88	64.65
Kelley-Inc	171	247	7.48	.44	9.46
bundle-Inc	196	162	58.21	9.11	68.85
BUNDLE+K-Inc	195	162	57.80	9.24	68.65

Table 4: CSP, minimization of the waste: averages on 20 random instances

We now come to the bin-packing problem, which is a special case of the cutting-stock problem where item demands take value 1. However, when an instance involves identical items (items with the same width) their demands are aggregated. Hence, bin-packing instances can be viewed as cutting-stock problems with low average demand (and typically more items). Our test problems are extracted from the OR library [3]. We use 20 random instances with 120 and 250 items, whose sizes are integer and vary between 20 to 100, while the size of the bin is 150. We use also 20 so-called triplets instances with 120 and 249 items, for which the bin capacity is 100, and the items sizes are real and range from 25 to 49.9. The triplets instances are generated from an optimum solution where there are exactly three items in each bin, which fill the bin capacity exactly (hence $\hat{u} = w/W$ is an optimal dual solution to (34)). They are known to be hard to solve for standard column generation algorithms. The results are given in Table 5. Bundle typically takes fewer iterations. It is also the fastest method on the hardest instances.

	Master	Counters		Timers (in seconds)		
		iter	cols	oracle	master	total
BP-120	KBASIC	222	223	0.79	.31	1.79
	KRICH	101	101	0.68	.19	1.15
	BUNDLE	107	90	.79	.23	1.36
BP-250	KBASIC	260	260	1.66	.50	2.90
	KRICH	134	134	1.35	.30	2.03
	BUNDLE	112	105	0.82	.30	1.60
BP-t120	KBASIC	339	340	5.82	.74	7.83
	KRICH	120	120	3.38	.34	4.14
	BUNDLE	80	79	2.17	.20	2.84
BP-t249	KBASIC	553	554	56.56	2.37	61.36
	KRICH	189	189	14.79	.81	16.49
	BUNDLE	139	138	7.20	.69	9.09

Table 5: Bin packing: results are averages on 20 instances

2.3 The vertex coloring problem (VCP)

Given a graph $G(V, E)$, the coloring problem is to use the minimum number of colors in assigning a color to each vertex so that no adjacent vertices receive the same color. This minimum number of colors $\chi(G)$ is called the *chromatic number*. Let N be an upper bound on $\chi(G)$. To formulate the coloring problem as (1), let us introduce variables $y^\ell \in \{0, 1\}$ to indicate whether color ℓ is used and $z_j^\ell \in \{0, 1\}$ to indicate whether vertex j is assigned color ℓ . Then $\chi(G)$ is the optimal value of

$$\min_{y, z \in \{0, 1\}} \sum_{\ell=1}^N y^\ell, \quad \sum_{\ell=1}^N z_j^\ell \geq 1, \quad j \in V, \quad z_i^\ell + z_j^\ell - y^\ell \leq 0, \quad (ij) \in E, \quad \ell = 1, \dots, N.$$

We dualize the covering constraints $\sum_{\ell=1}^N z_j^\ell \geq 1$ and keep the other constraints in X . The Lagrangian (8) is then

$$L(y, z, u) = \sum_{\ell=1}^N \left(y^\ell - \sum_{j \in V} u_j z_j^\ell \right) + \sum_{j \in V} u_j,$$

resulting in the dual function

$$\theta(u) = \sum_j u_j + \min_{y, z \in \{0, 1\}} \sum_{\ell} \left\{ y^\ell - \sum_j u_j z_j^\ell : z_i^\ell + z_j^\ell - y^\ell \leq 0, \quad (ij) \in E \right\}.$$

The above minimization problem is decomposable with respect to colors ℓ ; in fact, θ is obtained as the sum of N identical minimization problems. The value $y^\ell = 0$ produces $z^\ell = 0$, while for $y^\ell = 1$, $z = z^\ell$ solves a maximum independent set problem on G with weights u :

$$\max_{z \in \{0, 1\}} \sum_j u_j z_j, \quad z_i + z_j \leq 1, \quad (ij) \in E.$$

Denoting by $z(u)$ an optimal solution to this problem, the dual function is just as for CSP:

$$\theta(u) = 1u + N \min\{0, 1 - uz(u)\}.$$

Now, let us write the Dantzig-Wolfe master program (3): $\min \sum_{i \in I} y^i \lambda_i$, $\sum_{i \in I} z_j^i = 1$, $j \in V$, $\sum_{i \in I} \lambda \leq N$, $\lambda_i \geq 0$, $i \in I$, where I enumerates the solutions $(y, z) \in X^\ell$. As for CSP, $\sum_{i \in I} \lambda \leq N$ are inherited from original convexity constraints. Moreover, the constraint can be dropped given the objective. Thus the useful formulation is

$$\min \left\{ \sum_{i \in I} \lambda_i : \sum_{i \in I} z_j^i \geq 1, \quad j \in V, \quad \lambda_i \geq 0, \quad i \in I \right\},$$

Then, I enumerates independent sets of the graph G and x^i 's are indicator vectors of these sets. The dual is :

$$\max\{1u : uz^i \leq 1, i \in I, u \geq 0\}.$$

The column generation method will produce the so-called *fractional chromatic number*, denoted by $\chi_F(G)$, which is also an upper bound for the maximum clique number ω .

Numerical tests performed on two academic families of graphs (the Mycielski and Queen graphs), as well as on real-life instances (register allocation). These graphs were used as benchmarks for the DIMACS challenge³. Three methods are compared: KBASIC, KRICH and BUNDLE. The initial \hat{u} for BUNDLE is set to $\hat{u}_j = \text{deg}(j)/(|V| - 1)$ where $\text{deg}(j)$ is the number of neighbours of node j : the dual cost of assigning a color to a vertex connected to every other vertex is 1. KRICH is initialized with artificial variable cost set to $\alpha\hat{u}$ with $\alpha = 1.1$. The (NP-Hard) oracle subproblems are solved using a recursive enumeration algorithm [34].

The Mycielski graphs form a sequence of increasingly large graphs with clique number 2 (triangle-free graphs) and increasing chromatic number (for instance, Mycielski 2 is a pentagon). Since the fractional chromatic number is sandwiched between those two numbers, a large gap means a graph hard to color. Table 6 gives the computational results for these instances. Its first column gives, for each problem:

- the name of the problem,
- (number of nodes , number of edges),
- the clique, chromatic and fractional chromatic numbers: $\omega - \chi - \chi_F$.

Compared to KBASIC, the rough initialization of KRICH helps to reduce the computational time and the number of iterations. The computing time for the bundle method is competitive with KRICH, with a gain in solving the master.

Problem data	Master	Counters		Timers (in seconds)		
		iter	cols	oracle	master	total
Mycielski2 (5,5) 2 - 3 - 2.5	KBASIC	8	8	0.01	0.00	0.02
	KRICH	6	5	0.00	0.01	0.03
	BUNDLE	6	5	0.00	0.00	0.05
Mycielski3 (11,20) 2 - 4 - 2.9	KBASIC	29	29	0.01	0.01	0.09
	KRICH	12	11	0.00	0.00	0.07
	BUNDLE	19	11	0.04	0.00	0.09
Mycielski4 (23,71) 2 - 5 - 3.24	KBASIC	91	91	0.17	0.12	0.56
	KRICH	47	46	0.13	0.03	0.41
	BUNDLE	61	25	0.26	0.06	0.46
Mycielski5 (47,236) 2 - 6 - 3.55	KBASIC	264	264	2.79	0.74	4.96
	KRICH	177	176	2.83	0.69	5.05
	BUNDLE	177	55	4.59	0.53	5.76
Mycielski6 (95,755) 2 - 7 - 3.83	KBASIC	928	928	109.82	14.19	138.84
	KRICH	690	690	76.42	8.34	95.78
	BUNDLE	447	113	90.86	5.11	101.30
Mycielski6 (191,2360) 2 - 8 - 4.10	KBASIC	3724	3724	8875.99	997.05	10109.26
	KRICH	3204	3203	6704.52	684.32	7587.72
	BUNDLE	1254	510	7160.93	28.46	7217.13

Table 6: Vertex Coloring: results on Mycielski graphs

The Queen graphs arise from the puzzle of placing as many queens as possible on an $n \times n$ chessboard so that no two queens are on the same row, column or diagonal : graphs are constructed so that the answer to the puzzle correspond to a maximum independent set on the graph. The results on these highly symmetric instances are presented in Table 7, which reads as Table 6. Once again, stabilisation decreases substantially the number of iterations; but observe the corresponding increase in oracle's computing time, which can be dramatic.

In Table 8, we compare the three methods on 50-vertex subgraphs extracted from real instances arising in register allocation: assign variables of a program to high-speed access memory registers. On these graphs,

³<http://mat.gsia.cmu.edu/COLOR/instances.html>

Problem data	Master	Counters		Timers (in seconds)		
		iter	cols	oracle	master	total
Queen 5x5 (25,160) 5 – 5 – 5	KBASIC	60	61	0.14	0.03	0.25
	KRICH	6	6	0.01	0.01	0.08
	BUNDLE	11	7	0.05	0.01	0.13
Queen 6x6 (36,290) 6 – 7 – 7	KBASIC	84	84	0.25	0.12	0.59
	KRICH	53	53	0.27	0.18	0.70
	BUNDLE	44	41	0.62	0.07	0.85
Queen 7x7 (49,476) 7 – 7 – 7	KBASIC	141	142	1.15	0.32	1.86
	KRICH	206	206	3.65	1.11	5.83
	BUNDLE	14	11	1.08	0.05	1.25
Queen 8x8 (64,728) 8 – 9 – 8.44	KBASIC	197	198	10.23	0.67	11.64
	KRICH	134	134	8.83	0.40	9.96
	BUNDLE	74	70	30.98	0.28	31.82
Queen 9x9 (81,2112) 9 – 9 – 9	KBASIC	351	351	117.26	2.13	121.26
	KRICH	293	292	142.40	1.80	145.88
	BUNDLE	102	96	334.76	0.55	336.37
Queen 10x10 (100,2940) 10 – 10 – 10	KBASIC	457	458	1177.39	4.60	1185.11
	KRICH	349	349	1063.42	3.65	1069.64
	BUNDLE	312	156	6759.60	2.26	6764.44
Queen 11x11 (121,3960) 11 – 11 – 11	KBASIC	510	511	5912.88	8.01	5925.31
	KRICH	374	374	5457.54	6.07	5467.19
	BUNDLE	125	124	13017.26	3.88	13025.21

Table 7: Vertex Coloring: results on Queen graphs

computing times are much shorter using KBASIC. For KRICH, there are typically more vertices with strictly positive reward to consider when solving the subproblem, implying even more computing time than for BUNDLE. Moreover, the initial \hat{u} does not reflect the structure of the optimal dual solution, which typically involves three classes of vertices: those that receive weight 1 (as those in the largest clique), those with zero weight, and the “undecided”. Thus, several phases are required to eliminate all artificial columns from the solution.

2.4 The Capacitated Vehicle Routing Problem (CVRP)

Given a set of m customers, indexed by $j = 1, \dots, m$, demanding a delivery d_j , and N identical vehicles, indexed by $\ell = 1, \dots, N$, available to make these deliveries, the problem is to determine a route for each vehicle satisfying the following constraints:

- the route starts and ends at the depot (denoted by index 0);
- the deliveries assigned to a vehicle cannot exceed its capacity C ;
- each customer must be visited by some vehicle.

The distance matrix $\{c_{ij}\}_{0 \leq i, j \leq m}$ between customer sites is known (the distances are assumed to be integer and triangular inequalities hold). The goal is to minimize the sum of the lengths of the N routes.

Let $y_j^\ell = 1$ if customer j is visited by vehicle ℓ ; let $z_{ij}^\ell = 1$ if arc (i, j) is in the route of vehicle ℓ and zero otherwise; call q_j^ℓ the cumulated load of vehicle ℓ on departure from customer j . Then a possible formulation for (1) is:

$$\left\{ \begin{array}{ll} \min \sum_{i,j,\ell} c_{ij} z_{ij}^\ell & \\ \sum_{\ell=1}^N y_j^\ell \geq 1 & \text{for } j > 0 \\ \sum_i z_{ij}^\ell = \sum_i z_{ji}^\ell & \text{for all } j, \ell \\ \sum_i z_{ij}^\ell = y_j^\ell & \text{for all } j, \ell \\ q_i^\ell - q_j^\ell + C z_{ij}^\ell + d_j y_j^\ell \leq C & \text{for } i \neq j > 0 \text{ and all } \ell \\ d_j y_j^\ell \leq q_j^\ell \leq C y_j^\ell & \text{for } j > 0 \text{ and all } \ell \\ q_i^\ell \geq 0 & \text{for } j > 0 \text{ and all } \ell \\ y_j^\ell, z_{ij}^\ell \in \{0, 1\} & \text{for } i \neq j > 0 \text{ and all } \ell \end{array} \right. \quad (*)$$

Register instances	counters		timers (in seconds)		
	iter	cols	oracle	master	total
fpsol2.i.1	88	89	0.22	0.14	0.63
fpsol2.i.2	76	77	0.16	0.09	0.48
fpsol2.i.3	76	77	0.21	0.06	0.51
mulsol.i.1	125	126	0.40	0.16	0.90
mulsol.i.2	93	93	0.23	0.14	0.67
mulsol.i.3	93	93	0.24	0.11	0.61
mulsol.i.4	93	93	0.24	0.07	0.60
mulsol.i.5	93	93	0.25	0.10	0.63
zeroin.i.1	75	75	0.19	0.09	0.55
zeroin.i.2	75	75	0.14	0.10	0.49
zeroin.i.3	75	75	0.16	0.08	0.51
KBASIC average	87.5	87.8	0.22	0.10	0.60
fpsol2.i.1	117	106	3.31	0.21	3.78
fpsol2.i.2	142	114	7.85	0.26	8.50
fpsol2.i.3	142	114	7.88	0.22	8.38
mulsol.i.1	138	128	1.33	0.26	1.84
mulsol.i.2	276	247	2.13	0.45	3.28
mulsol.i.3	276	247	2.27	0.49	3.45
mulsol.i.4	276	247	2.25	0.51	3.43
mulsol.i.5	276	247	2.15	0.59	3.41
zeroin.i.1	146	134	4.89	0.25	5.50
zeroin.i.2	146	134	4.90	0.27	5.49
zeroin.i.3	146	134	4.85	0.22	5.42
KRICH average	189.2	168.4	3.98	0.34	4.77
fpsol2.i.1	62	58	0.61	0.09	0.92
fpsol2.i.2	71	62	0.62	0.15	1.00
fpsol2.i.3	71	62	0.68	0.19	1.03
mulsol.i.1	58	52	0.53	0.13	0.84
mulsol.i.2	51	40	0.41	0.11	0.68
mulsol.i.3	51	40	0.50	0.06	0.70
mulsol.i.4	51	40	0.46	0.11	0.70
mulsol.i.5	51	40	0.43	0.08	0.70
zeroin.i.1	64	54	0.60	0.14	0.93
zeroin.i.2	64	54	0.61	0.10	0.90
zeroin.i.3	64	54	0.59	0.13	0.90
BUNDLE average	59.8	50.5	0.55	0.12	0.85

Table 8: Vertex Coloring: results on register-allocation problems

Introducing the variable vector $x = (q, y, z)$, call $Ax \geq b$ the cover⁴ constraints (*); the other constraints are put in X . This gives the Lagrangian dual function (9)

$$\theta(u) = \sum_{j=1}^m u_j + \min_{(q,y,z) \in X} \sum_{\ell=1}^N \left(\sum_{i,j} c_{ij} z_{ij}^{\ell} - \sum_j u_j y_j^{\ell} \right)$$

which is decomposable with respect to ℓ : X is the Cartesian product of N identical subsets X^{ℓ} . The optimization sub-problem, $\min_{(q,y,z) \in X^{\ell}} \sum_{i,j} c_{ij} z_{ij} - \sum_j u_j y_j$ (where the useless index ℓ is dropped), is an elementary shortest path problem with resource constraints; it can be solved by dynamic programming [8]. Denoting its optimal solution by $x(u) = (q(u), y(u), z(u))$, the dual function has value:

$$\theta(u) = 1u + N[cz(u) - uy(u)]$$

An alternative presentation of the Lagrangian dual problem is Dantzig-Wolfe formulation (3). Here the master is:

$$\min \left\{ \sum_{i \in I} \left(\sum_{k,j} c_{kj} z_{kj}^i \right) \lambda_i : \sum_{i \in I} y_j \lambda_i \geq 1 \quad \forall j, \quad \sum_{i \in I} \lambda_i \leq N, \quad \lambda_i \geq 0 \quad i \in I \right\}$$

where I enumerates solutions $x^i = (q^i, y^i, z^i)$ of X^{ℓ} , i.e. feasible vehicle routes. Its dual takes the form:

$$\max \{ 1u - Nr : uy^i - r \leq cz^i, \quad i \in I \quad (u, r) \in \mathbb{R}_+^m \times \mathbb{R}_+ \}.$$

Kelley and bundle are compared on 12 test problems from the website [2]. The three P-instances and the E-instance are just those of [2]. The 8 remaining ones are reduced B-instances of [2], obtained by selecting a subset of N routes from the optimal solution and restricting the problem to the associated customers. The sizes m and N are provided in the problem names.

To define artificial column costs, we take an estimate edge cost \bar{c}_i for connecting client i , and then use

$$B := N \frac{\sum_j (c_{0j} + c_{j0})}{m} + \frac{m - N}{m} \sum_i \bar{c}_i,$$

which estimates the optimal cost of an integer solution. Indeed, a solution can be seen as a TSP tour visiting all customers and making N detours by the depot; for the inter customer distance, we take the average value $\frac{\sum_i \bar{c}_i}{m}$; for the cost of the detour to the depot, we take the average value $\frac{\sum_j (c_{0j} + c_{j0})}{m}$. One could take \bar{c}_i to be the average edge cost on, say, the 4 closest neighbors. A finer estimate is to take into account that customer with high demand are harder to accommodate in a cluster and, as a result, can be paired up with customer that are further away. Hence, we define $\bar{c}_i = \frac{1}{\beta_i} \sum_{j=1}^{\beta_i} c_{ij}$ where the neighbors, j , are assumed to be sorted in non-decreasing order of the edge costs, $\beta_i = 2(1 + \lceil \frac{d_i}{\bar{d}} \rceil)$, and \bar{d} is the average demand.

Seven versions are tested:

- KBASIC, where the master is initialized with an artificial column of cost equal to the above estimation B .
- KRICH initialized with m artificial columns associated with each client i . Their cost represents the fraction of total cost estimate B that is attributed to client j , i.e.

$$\hat{u}_j = \frac{w_j}{\sum_{k=1}^m w_k} B$$

where the weight factors are defined as $w_i = \bar{c}_i + \frac{1}{1 + \frac{c_{0i}}{\bar{d}}}$ to reflect a larger portion of the cost to customer with high connection cost, high demand (and thus lower number of neighbors with whom to share the cost of travelling from and to the depot) and that are far from the depot. When artificial columns remain in the solution we multiply them by $\alpha = 1.2$.

- BUNDLE0, initialized at $\hat{u} = 0$.
- BUNDLE, initialized at \hat{u} defined as for KRICH.
- KRICHMC is KRICH with multi-column generation for each subproblem: we insert in the master all feasible routes (or up to 500) found by the dynamic programming solver for the subproblem (this is standard practice when solving VRP by column generation).

⁴Covering constraints may be used instead of partitioning because triangular inequalities guarantee that a customer will not be visited more than once in an optimal solution.

- Kelley-Inc is initialized with a primal heuristic solution obtained using Clarke and Wright’s iterative route merging procedure (the artificial column of KBASIC is also included in case the heuristic solution is not feasible: i.e. if it uses too many vehicles).
- bundle-Inc, initialized at \hat{u} defined as for KRICH with an initial bundle made of the columns of Clarke and Wright’s solution.

We also tested BUNDLE+K but the results hardly differ from the use of bundle alone.

Averages	iter	cols	oracle	master	total
KBASIC	41	41	3.82	.10	4.10
KRICH	17	14	.74	.08	.91
KRICHMC	7	136	.85	.07	1.21
BUNDLE0	25	18	86.14	.05	86.32
BUNDLE	24	13	1.58	.06	1.73
Kelley-Inc	31	36	1.40	.09	1.66
bundle-Inc	24	18	1.58	.06	1.75

Table 9: Average results for CVRP instances

Problem	Method	Counters		Timers (seconds)		
		iter	cols	oracle	master	total
P-m16-N8	KBASIC	29	29	.15	.06	.31
	KRICH	18	17	.04	.06	.19
	BUNDLE0	30	22	.08	.06	.24
	BUNDLE	40	20	.18	.10	.42
	KRICHMC	6	42	.07	.03	.18
P-m22-N8	KBASIC	57	57	.54	.18	.90
	KRICH	37	31	.38	.17	.70
	BUNDLE0	46	34	.72	.16	1.10
	BUNDLE	38	26	.65	.11	.93
	KRICHMC	17	125	.61	.12	1.06
P-m23-N8	KBASIC	65	62	.72	.26	1.21
	KRICH	46	39	.70	.53	1.42
	BUNDLE0	70	39	1.51	.25	2.11
	BUNDLE	69	31	1.67	.24	2.13
	KRICHMC	22	145	1.13	.47	.02
E-m22-N4	KBASIC	83	83	15.41	.32	16.14
	KRICH	37	35	3.37	.11	3.71
	BUNDLE0	35	30	745.39	.05	745.65
	BUNDLE	44	23	5.58	.14	5.85
	KRICHMC	10	231	1.90	.10	2.78

Table 10: Detailed comparison on CVRP

Table 9 gives the average results. Details are given in Table 10 (for the first 4 instances) and Table 11 (the 8 remaining ones). These latter two tables explain in particular BUNDLE0’s large average computing time. Version KRICH is the fastest while KRICHMC needs fewer iterations. However, these results are sensitive to the fine tuning of the cost estimate and resulting value of \hat{u} (we previously tested rough estimates that yielded more iterations for KRICH). By contrast, bundle is quite robust: the number of iterations does not suffer from a poor initialization.

Most of the computing time is spent in the oracles. The resource-constrained shortest path problem solver of [8] enumerates many partial solutions when the rewards u_j are high and many clients seem attractive to include in the route. We noted experimentally that our initial \hat{u} overestimates on average the optimal dual solution: for most customers the dual values decrease in the course of the column generation procedure. Hence,

Problem	Method	Counters		Timers (seconds)		
		iter	cols	oracle	master	total
gen-B-m8-N2	KBASIC	21	21	.07	.04	.17
	KRICH	9	8	.06	.03	.12
	BUNDLE0	16	10	.13	.04	.22
	BUNDLE	12	7	.10	.01	.16
	KRICHMC	6	73	.27	.04	.44
gen-B-m12-N2	KBASIC	40	40	3.43	.09	3.72
	KRICH	12	11	1.45	.04	1.54
	BUNDLE0	23	20	748	.03	763
	BUNDLE	25	12	3.55	.03	3.68
	KRICHMC	7	282	2	.01	2.68
gen-B-m12-N2	KBASIC	34	34	2.57	.03	2.74
	KRICH	5	3	.15	.01	.19
	BUNDLE0	15	12	3.78	.02	3.86
	BUNDLE	10	4	.43	.01	.49
	KRICHMC	4	81	.31	.02	.41
gen-B-m14-N2	KBASIC	40	40	6.37	.05	6.67
	KRICH	4	3	.19	0	.26
	BUNDLE0	9	6	215.47	0	215.51
	BUNDLE	6	3	.45	0	.52
	KRICHMC	2	71	.23	0	.30
gen-B-m10-N2	KBASIC	29	29	.61	.06	.76
	KRICH	6	5	.21	.01	.26
	BUNDLE0	9	6	.75	.02	.80
	BUNDLE	14	7	1.52	.03	1.60
	KRICHMC	4	125	.69	.02	.83
gen-B-m13-N2	KBASIC	46	46	15.68	.09	16.03
	KRICH	6	5	1.79	.01	1.85
	BUNDLE0	17	14	57.29	.02	57.42
	BUNDLE	5	2	3.52	.01	3.57
	KRICHMC	4	224	2.02	.01	2.33
gen-B-m10-N2	KBASIC	27	28	.22	.04	.35
	KRICH	13	12	.32	.06	.42
	BUNDLE0	21	17	.69	.03	.80
	BUNDLE	23	13	.95	.03	1.03
	KRICHMC	6	138	.45	.02	.72
gen-B-m10-N2	KBASIC	25	25	.15	.01	.29
	KRICH	13	10	.24	.03	.30
	BUNDLE0	19	14	.42	.03	.53
	BUNDLE	13	8	.40	.01	.44
	KRICHMC	7	106	.55	0	.80

Table 11: Detailed comparison on reduced CVRP

the most expensive oracle calls tend to be at the beginning of the column generation procedure. For KBASIC and BUNDLE0 however, many rewards u_j are very small in the initial iterations (initial extreme dual solutions under KBASIC concentrate all the reward on a few clients) and the oracle preprocessor removes the unattractive nodes, making the problem easier. After a few timid iterations, BUNDLE0 increases the u 's largely and uniformly, which yields expensive oracles. This behaviour is illustrated by Fig. 7.




Figure 7: CVRP: oracle computing time for instance E-m22-N4

2.5 The multi-item lot-sizing (MILS) problem

The problem is to optimize the production planning of items $\ell = 1, \dots, N$ for time periods $j = 1, \dots, m$. Before production can take place, the machine must be specifically setup for the item (at a cost f_j^ℓ and with a loss of production capacity b_j^ℓ). The other costs associated with item ℓ are unit production costs p_j^ℓ and holding costs h_j^ℓ . Item demands for each period j are denoted by d_j^ℓ and the machine production capacity in period j is C_j . Set the binary variable y_j^ℓ to 1 if the machine is setup for item ℓ in period j . Let z_j^ℓ be the production of item ℓ in period j and s_j^ℓ the stock of ℓ at the end of period j .

Then (1) takes the form:

$$\begin{cases} \min \sum_{j,\ell} (f_j^\ell y_j^\ell + p_j^\ell z_j^\ell + h_j^\ell s_j^\ell) & \text{for all } j & (*) \\ \sum_{\ell=1}^N z_j^\ell + b_j^\ell y_j^\ell \leq C_j & \text{for all } j, \ell \\ s_{j-1}^\ell + z_j^\ell = d_j^\ell + s_j^\ell & \text{for all } j, \ell \\ z_j^\ell \leq D_j^\ell y_j^\ell & \text{for all } j, \ell \\ z_j^\ell, s_j^\ell \geq 0 & \text{for all } j, \ell \\ y_j^\ell \in \{0, 1\} & \text{for all } j, \ell \end{cases}$$

Here $D_j^\ell = \sum_{\tau=j}^m d_\tau^\ell$; the corresponding constraint forces $z_j^\ell = 0$ unless $y_j^\ell = 1$.

Introducing the variable vector $x = (y, z, s)$, call $Ax \geq b$ the capacity constraints (*) and put the other constraints in X ; this gives the Lagrangian dual function (9)

$$\theta(u) = - \sum_j u_j C_j + \min_{(y,z,s) \in X} \sum_j [(f_j^\ell + b_j^\ell u_j) y_j^\ell + (p_j^\ell + u_j) z_j^\ell + h_j^\ell s_j^\ell]$$

which is decomposable with respect to ℓ : X is the Cartesian product of N non-identical subsets X^ℓ . Each optimization sub-problem

$$\theta^\ell(u) := \min_{(y,z,s) \in X^\ell} \sum_t [(f_j + b_j u_j) y_j + (p_j + u_j) z_j + h_j s_j] \quad (36)$$

(where the useless index ℓ is dropped) is a single-item lot-sizing problem with unbounded capacity; it can be solved by dynamic programming in polynomial time. The dual function is then

$$\theta(u) = \sum_{\ell=1}^N \theta^\ell(u) - \sum_{j=1}^m C_j u_j.$$

The master formulation (3) takes the form:

$$\min \left\{ \sum_{\ell, i \in I^\ell} \left(\sum_j (f_j^\ell y_j^i + p_j^\ell z_j^i + h_j^\ell s_j^i) \right) \lambda_i^\ell \right. \\ \left. \sum_{\ell, i \in I^\ell} (z_j^i + b_j^\ell y_j^i) \lambda_i^\ell \leq C_j \quad \forall j, \right. \quad (37)$$

$$\sum_{i \in I^\ell} \lambda_i^\ell \geq 1 \quad \forall \ell, \quad (38) \\ \lambda_i^\ell \geq 0 \quad \forall \ell, i \in I^\ell \}$$

where I^ℓ enumerates production plans $x^i = (y^i, z^i, s^i)$ of X^ℓ , i.e. $X^\ell = \{x^i\}_{i \in I^\ell}$. Its dual takes the form:

$$\max \left\{ \sum_{\ell=1}^N r_\ell - \sum_j u_j C_j : \right. \\ \left. r_\ell - \sum_j u_j (z_j^\ell + b_j^\ell y_j^\ell) \leq (f_j^\ell y_j^\ell + p_j^\ell z_j^\ell + h_j^\ell s_j^\ell) \quad \forall \ell, i \in I^\ell, \right. \\ \left. (u, r) \in \mathbb{R}_+^m \times \mathbb{R}_+^N \right\}.$$

This model has an important implementation feature: in view of (36), the dual function θ is a sum of N *different* concave functions θ^ℓ . The approximation θ^k of (12) at iteration k can therefore be *disaggregated* in the sum of N “restricted local” dual functions

$$\theta_k^\ell(u) := \min_{i \in I_k^\ell} L^\ell(x^i, u);$$

here $L^\ell(x^i, u)$ denotes the Lagrangian corresponding to product ℓ . Such a disaggregation results in a more accurate restricted dual function. In our experiments, Kelley uses disaggregation; but, even though this would be possible for bundle (as described and assessed for example in [31]), we have not implemented the corresponding software. It should be mentioned that the use of disaggregation in bundle would probably result in fewer iterations and more CPU time in the master.

In Table 12, we compare Kelley and bundle on randomly generated instances with $N = 20$ items and $m = 60$ periods. Each line is an average over 10 instances. For version KBASIC, the master is initialized with N artificial columns – one for each subproblem – with cost equal to a large constant: the artificial column associated with item ℓ has coefficient 1 in the associated item covering constraint (38) and zero elsewhere. For version KRICH, the master is initialized with $m + N$ artificial columns – one for each capacity constraint (37) and one for each item covering constraint (38) – with coefficient 1 in the associated constraint and zero elsewhere. Their cost is an estimation of the dual value associated with the constraint: for \hat{r}_ℓ , we use the cost of the production planning solving the single-item lot-sizing problem with unbounded capacity for item ℓ ; $\sum_\ell \hat{r}_\ell$ defines a lower bound on the multi-item lot-sizing problem; we then arbitrarily assume that the capacity constraints will yield a 20% cost increase over this bound and we define $\hat{u}_j = \frac{0.2 \sum_\ell \hat{r}_\ell}{m C_j}$; the artificial variable costs are set to $\alpha \hat{u}_j$ and $\alpha \hat{r}_\ell$ respectively, with $\alpha = 1.2$. The above \hat{u} is used to initialize BUNDLE, while the initial \hat{u} is 0 for BUNDLE0. Counter “Sp” indicates the total number of subproblems solved; the reported number “col” of different columns is likewise disaggregated (for KBASIC, cols includes $N = 20$ artificial columns).

Master	Counters			Timers (seconds)		
	iter	Sp	cols	oracle	master	total
KBASIC	9	186	118	0.72	0.01	2.63
KRICH	9	180	90	0.67	0.01	2.50
BUNDLE0	37	754	80	3.94	0.06	5.64
BUNDLE	53	1062	303	5.59	0.07	9.95

Table 12: MILS: average results on 10 random instances with 20 items and 60 periods

2.6 The traveling salesman problem (TSP)

Given a complete undirected graph $G = (V, E)$ and length c_e for each edge $e \in E$, the problem is to find a tour of minimum length. For $S \subset V$, let $\delta(S)$ be the set of edges $\{i, j\}$ with $i \in S$ and $j \notin S$, and let $E(S)$ be the set of edges $\{i, j\}$ with $i \in S$ and $j \in S$. To make our notation consistent with the rest of the paper, we set $V = \{0, 1, \dots, m\}$; a possible formulation for (1) is

$$\begin{cases} \min \sum_{e \in E} c_e x_e, \\ \sum_{e \in \delta(\{j\})} x_e = 2 & \text{for } j = 0, 1, \dots, m, \\ \sum_{e \in E(S)} x_e \leq |S| - 1 & \text{for all } \emptyset \neq S \subseteq \{1, \dots, m\}, \\ \sum_{e \notin \delta(\{0\})} x_e = n - 2, \\ x_e \in \{0, 1\} & \text{for } e \in E. \end{cases} \quad (*)$$

Call $Ax \geq b$ those constraints (*) with $j > 0$ (they can in fact be replaced by inequalities) and put the other constraints in X : the dual variables are u_1, \dots, u_m . For notational convenience, we set $u_0 = 0$, and the Lagrangian is

$$L(x, u) = \sum_{e=\{i,j\} \in E} (c_e - u_i - u_j)x_e + 2 \sum_{j=1}^m u_j.$$

This relaxation was introduced by M. Held and R. Karp: [14, 15].

In the definition of X , no two variables x_e and $x_{e'}$ appear in the same constraint if $e \in \delta(\{0\})$ and $e' \notin \delta(\{0\})$. As a result, the solutions of the oracle are the minimum-weight 1-trees, i.e. the minimum-weight spanning trees on $\{1, \dots, m\}$ completed by the two minimum-weight edges of $\delta(\{0\})$. The oracle is then an easy problem which can be solved by a Prim's algorithm [41].

Let I enumerate the 1-trees, i.e. the solutions $x^i \in X$. Then, the master program (3) of a Dantzig-Wolfe reformulation is

$$\begin{cases} \min \sum_{i \in I} \lambda_i \sum_{e \in E} c_e x_e^i, \\ \sum_{i \in I} \lambda_i \sum_{e \in \delta(\{j\})} x_e^i = 2 & \text{for all } j \in V \setminus \{0\}, \\ \sum_{i \in I} \lambda_i \leq 1, \\ \lambda_i \geq 0 & \text{for all } i \in I. \end{cases}$$

Its dual (11) takes the form

$$\begin{cases} \max 2u - r & (u, r) \in \mathbb{R}_+^m, \\ \sum_{j \in V \setminus \{0\}} u_j \sum_{e \in \delta(\{j\})} x_e^i - r \leq \sum_{e \in E} c_e x_e^i & \text{for all } i \in I. \end{cases}$$




Figure 8: TSP: results with bays29

We have applied KBASIC, KRICH and BUNDLE0 on three instances from tsplib⁵ having 29, 76 and 442 cities respectively. For KRICH, we compute

$$\hat{u}_j := \min \left\{ \sum_{e \in \delta(\{j\})} c_e x_e : \sum_{e \in \delta(\{j\})} x_e = 2, x_e \in \{0, 1\} \text{ for } e \in \delta(\{j\}) \right\};$$

then we define m artificial columns with cost $\alpha \hat{u}$ for $\alpha = 1.2$. The results are reported in Figures 8 to 10. On the third example (pcb442), no convergence can be obtained from Kelley, even with a rich initialization. The reported primal-dual gap is the value $c\hat{x} - \theta(u)$, but with a \hat{x} that may still involve artificial columns (then $c\hat{x}$ is not a true primal bound).

Finally, Fig. 11 displays the 1-trees computed by the oracle during the first 39 [resp. 31] iterations of Kelley [resp. BUNDLE0].

⁵<http://www.iwr.uni-heidelberg.de/groups/comopt/software/TSPLIB95/>




Figure 9: TSP: results with pr76




Figure 10: TSP: results with pcb442




Figure 11: TSP: primal behaviour on Bays29 for KBASIC, KRICH (middle) and BUNDLE0 (bottom)

Conclusions

The first part of this paper has presented in a unified way various stabilizing schemes of column generation, including the bundle and ACCPM approaches. Stabilization is conveniently introduced in the dual space; deriving its primal counterpart is easy when an LP formulation exists. However, we have mentioned a possible

dualization scheme allowing the same primal derivation, even when the stabilized master is a genuine nonlinear program.

The above-mentioned work offers little new material, as it uses only well-known tools from convex analysis; besides, it is also treated in [9]. However, the innovative part of this paper is really §2, where the numerical behaviour of the bundle method is compared with standard column generation on a few combinatorial problems. Some conclusions can be proposed from our experiments.

- (i) Sometimes Kelley and bundle are grossly comparable in terms of number of calls to the oracle. This is generally observed for problems of §2.2 to §2.5.

It is worth recalling here Remark 2.3: in the case of CSP and VCP, we use our bundle code to solve a constrained problem by exact penalty, which is known to be rather inefficient. Using a constrained bundle method such as [22, 30] would probably give better results.

Another useful observation is the lack of robustness of numerical experiments in nonsmooth optimization. Because the $x = x(u)$ answered by the oracle is discontinuous, small changes (in the data, the initialization, the accuracy of the master, ...) may result in substantial changes in the behaviour of the column generation process. To solve a given problem by a given method, it is common that *the same* software takes substantially more iterations (say 20%) when compiled on a different computer. Keeping this in mind, the differences reported in §2.2 to §2.5 should be taken with care.

- (ii) Sometimes, typically when m grows, Kelley simply collapses: see §2.6. This phenomenon was known before, the interesting point is that bundle still works for much larger instances: we refer to [28], in which instances with 10^3 vertices are solved to (dual) optimality; ACCPM is reported to behave similarly.
- (iii) An interesting question related with (ii) is: when should Kelley behave poorly compared to bundle? According to Nemirovskii's example (see the end of §1.1), this should be more likely for large m . Common sense suggests that it should generally happen when θ^k of (12) approximates poorly the true dual function θ of (9); such is the case when the graph of θ presents many small facets. Figure 12 gives a rudimentary illustration of this point: both functions θ on the left and on the right are piecewise linear; nevertheless, approximating the one on the left by a piecewise linear model θ^k requires many sampling points; a quadratic model should probably be more efficient – and this is precisely what bundle proposes.




Figure 12: The graph of the dual function

Anyway, it seems dary to propose a reliable answer to this question (assuming that a reliable answer exists at all!)

- (iv) A phenomenon, already known and confirmed by our experiments, is that stabilization may affect the computing time spent in the oracle. The choice of the oracle solver (a branch-and-bound versus a dynamic program, for instance), its implementation and even the initialization of the master program, are determinant factors when analyzing oracle computing time over the course of the column generation algorithm. Nevertheless, even though the reported times should be taken cautiously, our NP oracles may be perceivably more difficult when called by a stabilized method. Paroxystic entries such as Queen 11x11 in Table 7 or E-m22-N4 in Table 9 can hardly be tolerated. We have not fully explained this behaviour, which may occur at any time during the column generation process: either by the end (CSP) or at the beginning (CVRP).

On the other hand, the dual iterates u^k of a stabilized algorithm are supposedly close together. It might therefore be advantageous to use “warm-started” oracles *à la* [46].

- (v) Being specifically designed to maximize a concave function such as θ , the bundle method needs an *exact* oracle, to compute exact values of θ . Keeping in mind (iv) above, this can become a killing disadvantage; by contrast, Kelley simply needs separating hyperplanes (which do not have to touch the graph of θ). One of the consequences is that, for problems with NP oracle – especially VCP and CVRP –, our experiments were limited to relatively small instances.

The cure lies in “robust” stabilized algorithms, accepting *inaccurate* oracles which do not fully minimize the Lagrangian. Concerning the bundle method, a promising step forward is made along these lines in [24]: preliminary experiments show that the benefit can be drastic.

- (vi) Perhaps the most important message delivered by these experiments is that, in terms of CPU, the price to pay for solving the quadratic master (28) is quite reasonable with respect to the LP master (6). Here this price is negligible but the dual problems to solve are fairly small, and the oracles are expensive. Nevertheless, QP is also cheap for TSP, even large ones – see [28].

Besides, it should be observed that little R&D work has been devoted so far to QP, compared to the 50 years of intense activity in LP technology. The relative price of QP can therefore be expected to decrease in the future.

- (vii) Finally it can also be mentioned that an appropriate initialization of KRICH can require fairly delicate heuristic considerations. The paper provides such dual heuristics. They do yield significant reduction in the number of iterations, compared with poorer initialization on which we did not report. On the other hand, they can be sensitive to the characteristics of the datasets; VCP is an illustration. By contrast, bundle appears as a push-button method, in which even the initial \hat{u} has little importance.

References

- [1] K. Anstreicher and L.A. Wolsey. On dual solutions in subgradient optimization. Unpublished manuscript, CORE, Louvain-la-Neuve, Belgium, 1993.
- [2] P. Augerat. VRP problem instances, 1995. <http://www.branchandcut.org/VRP/data/>.
- [3] J.E. Beasley. Or-library: distributing test problems by electronic mail. *Journal of the Operational Research Society*, 41(11):1069–1072, 1990. <http://mscmga.ms.ic.ac.uk/jeb/orlib/binpackinfo.html>.
- [4] H. Ben-Amor and J. Desrosiers. A proximal-like algorithm for column generation stabilization. Technical Report G-2003-43, Les Cahiers du Gerad, Montréal, 2003.
- [5] E. Cheney and A. Goldstein. Newton’s method for convex programming and Tchebycheff approximations. *Numerische Mathematik*, 1:253–268, 1959.
- [6] Dash Optimization. *Xpress-MP: User guide and Reference Manual, Release 12*, 2001. <http://www.dashoptimization.com>.
- [7] Y.M. Ermol’ev. Methods of solution of nonlinear extremal problems. *Kibernetika*, 2(4):1–17, 1966.
- [8] D. Feillet, P. Dejax, M. Gendreau, and C. Gueguen. An exact algorithm for the elementary shortest path problem with resource constraints: application to some vehicle routing problems. *Networks*, 44(3):216–229, 2004.
- [9] A. Frangioni. Generalized bundle methods. *SIAM Journal on Optimization*, 13(1):117–156, 2002.
- [10] A.M. Geoffrion. Lagrangean relaxation for integer programming. *Mathematical Programming Study*, 2:82–114, 1974.
- [11] J.-L. Goffin, A. Haurie, and J.-Ph. Vial. Decomposition and nondifferentiable optimization with the projective algorithm. *Management Science*, 38(2):284–302, 1992.
- [12] J.-L. Goffin, Z.-Q. Luo, and Y. Ye. Complexity analysis of an interior cutting plane for convex feasibility problems. *SIAM Journal on Optimization*, 6(3):638–652, 1996.
- [13] J.-L. Goffin and J.-Ph. Vial. Convex nondifferentiable optimization: a survey focused on the analytic center cutting plane method. *Optimization Methods and Software*, 17(5):805–867, 2002.
- [14] M. Held and R. Karp. The traveling salesman problem and minimum spanning trees. *Operations Research*, 18:1138–1162, 1970.
- [15] M. Held and R. Karp. The traveling salesman problem and minimum spanning trees: Part II. *Mathematical Programming*, 1(1):6–25, 1971.

- [16] J.-B. Hiriart-Urruty and C. Lemaréchal. *Convex Analysis and Minimization Algorithms*. Springer Verlag, Heidelberg, 1993. Two volumes.
- [17] J.-B. Hiriart-Urruty and C. Lemaréchal. *Fundamentals of Convex Analysis*. Springer Verlag, Heidelberg, 2001.
- [18] J. E. Kelley. The cutting plane method for solving convex programs. *J. Soc. Indust. Appl. Math.*, 8:703–712, 1960.
- [19] S. Kim, K.N. Chang, and J.Y. Lee. A descent method with linear programming subproblems for nondifferentiable convex optimization. *Mathematical Programming*, 71(1):17–28, 1995.
- [20] K. C. Kiwiel. A dual method for certain positive semidefinite quadratic programming problems. *SIAM Journal on Scientific and Statistical Computing*, 10(1):175–186, 1989.
- [21] K.C. Kiwiel. An aggregate subgradient method for nonsmooth convex minimization. *Mathematical Programming*, 27:320–341, 1983.
- [22] K.C. Kiwiel. *Methods of Descent for Nondifferentiable Optimization*. Lecture Notes in Mathematics 1133. Springer Verlag, Heidelberg, 1985.
- [23] K.C. Kiwiel. A Cholesky dual method for proximal piecewise linear programming. *Numerische Mathematik*, 68:325–340, 1994.
- [24] K.C. Kiwiel. A proximal bundle method with approximate subgradient linearizations. Systems Research Institute, Warsaw; to appear in Siam J. on Optimization, 2002.
- [25] T. Larsson, M. Patriksson, and A.B. Strömberg. Ergodic, primal convergence in dual subgradient schemes for convex programming. *Mathematical Programming*, 86(2):283–312, 1999.
- [26] C. Lemaréchal. An algorithm for minimizing convex functions. In J.L. Rosenfeld, editor, *Information Processing '74*, pages 552–556. North Holland, 1974.
- [27] C. Lemaréchal. Nonsmooth optimization and descent methods. Research Report 78-4, IIASA, 1978.
- [28] C. Lemaréchal. Lagrangian relaxation. In M. Jünger and D. Naddef, editors, *Computational Combinatorial Optimization*, pages 115–160. Springer Verlag, Heidelberg, 2001.
- [29] C. Lemaréchal. The omnipresence of Lagrange. *4OR*, 1(1):7,25, 2003.
- [30] C. Lemaréchal, A.S. Nemirovskii, and Yu.E. Nesterov. New variants of bundle methods. *Mathematical Programming*, 69:111–148, 1995.
- [31] C. Lemaréchal, F. Pellegrino, A. Renaud, and C. Sagastizábal. Bundle methods applied to the unit-commitment problem. In J. Doležal and J. Fidler, editors, *System Modelling and Optimization*, pages 395–402. Chapman and Hall, 1996.
- [32] C. Lemaréchal and C. Sagastizábal. Variable metric bundle methods: from conceptual to implementable forms. *Mathematical Programming*, 76(3):393–410, 1997.
- [33] R.E. Marsten, W.W. Hogan, and J.W. Blankenship. The boxstep method for large-scale optimization. *Operations Research*, 23(3):389–405, 1975.
- [34] A. Mehrotra and M.A. Trick. A column generation approach to graph coloring. *INFORMS J. on Computing*, 8(4):344–354, 1996.
- [35] O. du Merle, D. Villeneuve, J. Desrosiers, and P. Hansen. Stabilized column generation. *Discrete Mathematics*, 194:229–237, 1999.
- [36] A.S. Nemirovskii and D. Yudin. Informational complexity and efficient methods for the solution of convex extremal problems. *Ékonomika i Matematicheskie Metody*, 12:357–369, 1976. (in Russian. English translation: *Matekon*, 13, 3-25).
- [37] A.S. Nemirovskii and D. Yudin. *Problem Complexity and Method Efficiency in Optimization*. Wiley-Interscience Series in Discrete Mathematics, 1983. (Original Russian: Nauka, 1979).

- [38] Yu.E. Nesterov. Complexity estimates of some cutting plane methods based on the analytic barrier. *Mathematical Programming*, 69(1):149–176, 1995.
- [39] Yu.E. Nesterov and J.-Ph. Vial. Homogeneous analytic center cutting plane methods for convex problems and variational inequalities. *SIAM Journal on Optimization*, 9(3):707–728, 1999.
- [40] B.T. Polyak. A general method for solving extremum problems. *Soviet Mathematics Doklady*, 8:593–597, 1967.
- [41] R. C. Prim. Shortest connection networks and some generalizations. *Bell System Technological Journal*, 36:1389–1401, 1957.
- [42] R.T. Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- [43] R.T. Rockafellar. A dual approach to solving nonlinear programming problems by constrained optimization. *Mathematical Programming*, 5:354–373, 1973.
- [44] N. Shor. Utilization of the operation of space dilatation in the minimization of convex functions. *Cybernetics*, 6(1):7–15, 1970.
- [45] N.Z. Shor. *Minimization methods for non-differentiable functions*. Springer Verlag, Berlin, 1985.
- [46] B. Thiongane, A. Nagih, and G. Plateau. Adapted step size in a 0-1 knapsack lagrangean dual solving algorithm. *Annals of Operations Research*, 2004.
- [47] H. Uzawa. Iterative methods for concave programming. In K. Arrow, L. Hurwicz, and H. Uzawa, editors, *Studies in Linear and Nonlinear Programming*, pages 154–165. Stanford University Press, 1959.
- [48] F. Vanderbeck. Extending Dantzig’s bound to the bounded multi-class binary knapsack problem. *Mathematical Programming*, 94(1):125–16, 2002.
- [49] F. Vanderbeck. Dantzig-Wolfe re-formulation or how to exploit simultaneously original formulation and column generation re-formulation. Working paper U-03.24, Univ. Bordeaux 1, Talence, France, 2003.



Unité de recherche INRIA Rhône-Alpes

655, avenue de l'Europe - 38330 Montbonnot-St-Martin (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique

615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur

INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)

<http://www.inria.fr>

ISSN 0249-6399