
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Boland, Natasha; Christiansen, Jeffrey ; Dandurand, Brian ; Eberhard, Andrew; Linderoth, Jeff; Luedtke, James; Oliveira, Fabricio

Combining Progressive Hedging with a Frank--Wolfe Method to Compute Lagrangian Dual Bounds in Stochastic Mixed-Integer Programming

Published in:
SIAM JOURNAL ON OPTIMIZATION

DOI:
[10.1137/16M1076290](https://doi.org/10.1137/16M1076290)

Published: 01/01/2018

Document Version
Publisher's PDF, also known as Version of record

Please cite the original version:
Boland, N., Christiansen, J., Dandurand, B., Eberhard, A., Linderoth, J., Luedtke, J., & Oliveira, F. (2018). Combining Progressive Hedging with a Frank--Wolfe Method to Compute Lagrangian Dual Bounds in Stochastic Mixed-Integer Programming. SIAM JOURNAL ON OPTIMIZATION, 28(2), 1312-1336.
<https://doi.org/10.1137/16M1076290>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

COMBINING PROGRESSIVE HEDGING WITH A FRANK–WOLFE METHOD TO COMPUTE LAGRANGIAN DUAL BOUNDS IN STOCHASTIC MIXED-INTEGER PROGRAMMING*

NATASHIA BOLAND[†], JEFFREY CHRISTIANSEN[‡], BRIAN DANDURAND[‡],
ANDREW EBERHARD[‡], JEFF LINDEROTH[§], JAMES LUEDTKE[§], AND
FABRICIO OLIVEIRA[‡]

Abstract. We present a new primal-dual algorithm for computing the value of the Lagrangian dual of a stochastic mixed-integer program (SMIP) formed by relaxing its nonanticipativity constraints. This dual is widely used in decomposition methods for the solution of SMIPs. The algorithm relies on the well-known progressive hedging method, but unlike previous progressive hedging approaches for SMIP, our algorithm can be shown to converge to the optimal Lagrangian dual value. The key improvement in the new algorithm is an inner loop of optimized linearization steps, similar to those taken in the classical Frank–Wolfe method. Numerical results demonstrate that our new algorithm empirically outperforms the standard implementation of progressive hedging for obtaining bounds in SMIP.

Key words. mixed-integer stochastic programming, Lagrangian duality, progressive hedging, Frank–Wolfe method

AMS subject classifications. 90C06, 90C11, 90C15, 90C46

DOI. 10.1137/16M1076290

1. Introduction. Stochastic programming with recourse provides a framework for modeling problems where decisions are made in stages. Between stages, some uncertainty in the problem parameters is unveiled, and decisions in subsequent stages may depend on the outcome of this uncertainty. When some decisions are modeled using discrete variables, the problem is known as a stochastic mixed-integer programming (SMIP) problem. The ability to simultaneously model uncertainty and discrete decisions makes SMIP a powerful modeling paradigm for applications. Important applications employing SMIP models include unit commitment and hydro-thermal generation scheduling [26, 36], military operations [34], vaccination planning [30, 37], air traffic flow management [4], forestry management and forest fire response [6, 28], and supply chain and logistics planning [20, 22]. However, the combination of uncertainty and discreteness makes this class of problems extremely challenging from a computational perspective. In this paper, we present a new and effective algorithm for computing lower bounds that arise from a Lagrangian-relaxation approach.

*Received by the editors May 20, 2016; accepted for publication (in revised form) January 17, 2018; published electronically May 8, 2018.

<http://www.siam.org/journals/siopt/28-2/M107629.html>

Funding: The work of authors Boland, Christiansen, Dandurand, Eberhard, Linderoth, and Oliveira was supported in part or in whole by the Australian Research Council (ARC) grant ARC DP140100985. The work of authors Linderoth and Luedtke was supported in part by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under contract number DE-AC02-06CH11357 and by the NSF under award 1634597.

[†]Georgia Institute of Technology, Atlanta, GA 30332 (natashia.boland@gmail.com).

[‡]RMIT University, Melbourne, Victoria, Australia (s3507717@student.rmit.edu.au, brian.dandurand@rmit.edu.au, andy.eberhard@rmit.edu.au, fabricio.oliveira@aalto.fi).

[§]Department of Industrial and Systems Engineering, Wisconsin Institutes of Discovery, University of Wisconsin-Madison, Madison, WI 53706 (linderoth@wisc.edu, jim.luedtke@wisc.edu).

The mathematical statement of a two-stage SMIP is

$$(1) \quad \zeta^{SMIP} := \min_x \{c^\top x + \mathcal{Q}(x) : x \in X\},$$

where the vector $c \in \mathbb{R}^{n_x}$ is known, and X is a mixed-integer linear set consisting of linear constraints and integer restrictions on some components of x . The function $\mathcal{Q} : \mathbb{R}^{n_x} \mapsto \mathbb{R}$ is the expected recourse value

$$\mathcal{Q}(x) := \mathbb{E}_\xi \left[\min_y \{q(\xi)^\top y : W(\xi)y = h(\xi) - T(\xi)x, y \in Y(\xi)\} \right].$$

We assume that the random variable ξ is taken from a discrete distribution indexed by the finite set \mathcal{S} , consisting of the realizations, $\xi_1, \dots, \xi_{|\mathcal{S}|}$, corresponding to strictly positive probabilities of realization, $p_1, \dots, p_{|\mathcal{S}|}$. When ξ is not discrete, a finite scenario approximation can be obtained via Monte Carlo sampling [19, 24] or other methods [11, 10]. Each realization ξ_s of ξ is called a *scenario* and encodes the realizations observed for each of the random elements $(q(\xi_s), h(\xi_s), W(\xi_s), T(\xi_s), Y(\xi_s))$. For notational brevity, we refer to this collection of random elements respectively as $(q_s, h_s, W_s, T_s, Y_s)$. For each $s \in \mathcal{S}$, the set $Y_s \subset \mathbb{R}^{n_y}$ is a mixed-integer set containing both linear constraints and integrality constraints on a subset of the variables y_s .

The problem (1) may be reformulated as a *deterministic equivalent*

$$(2) \quad \zeta^{SMIP} = \min_{x, y} \left\{ c^\top x + \sum_{s \in \mathcal{S}} p_s q_s^\top y_s : (x, y_s) \in K_s \ \forall s \in \mathcal{S} \right\},$$

where $K_s := \{(x, y_s) : W_s y_s = h_s - T_s x, x \in X, y_s \in Y_s\}$. Problem (2) has a special structure that can be algorithmically exploited by decomposition methods. To induce a decomposable structure, scenario-dependent copies x_s for each $s \in \mathcal{S}$ of the first-stage variable x are introduced to create the following reformulation of (2):

$$(3) \quad \zeta^{SMIP} = \min_{x, y, z} \left\{ \sum_{s \in \mathcal{S}} p_s (c^\top x_s + q_s^\top y_s) : (x_s, y_s) \in K_s, x_s = z \ \forall s \in \mathcal{S}, z \in \mathbb{R}^{n_x} \right\}.$$

The constraints $x_s = z, s \in \mathcal{S}$, enforce nonanticipativity for first-stage decisions; the first-stage decisions x_s must be the same (z) for each scenario $s \in \mathcal{S}$. Applying Lagrangian relaxation to the nonanticipativity constraints in problem (3) yields the *nonanticipative Lagrangian dual function*

$$(4) \quad \phi(\mu) := \min_{x, y, z} \left\{ \sum_{s \in \mathcal{S}} p_s [c^\top x_s + q_s^\top y_s + \mu_s^\top (x_s - z)] : (x_s, y_s) \in K_s \ \forall s \in \mathcal{S}, z \in \mathbb{R}^{n_x} \right\},$$

where $\mu = (\mu_1, \dots, \mu_{|\mathcal{S}|}) \in \prod_{s \in \mathcal{S}} \mathbb{R}^{n_x}$ is the vector of multipliers associated with the relaxed constraints $x_s = z, s \in \mathcal{S}$. By setting $\omega_s := \frac{1}{p_s} \mu_s$, (4) may be rewritten as

$$(5) \quad \phi(\omega) := \min_{x, y, z} \left\{ \sum_{s \in \mathcal{S}} p_s L_s(x_s, y_s, z, \omega_s) : (x_s, y_s) \in K_s \ \forall s \in \mathcal{S}, z \in \mathbb{R}^{n_x} \right\},$$

where

$$L_s(x_s, y_s, z, \omega_s) := c^\top x_s + q_s^\top y_s + \omega_s^\top (x_s - z).$$

Since z is unconstrained in the optimization problem in the definition (5), in order for the Lagrangian function $\phi(\omega)$ to be bounded from below, we require as a condition

of dual feasibility that $\sum_{s \in \mathcal{S}} p_s \omega_s = 0$. Under this assumption, the z term vanishes, and the Lagrangian dual function (5) decomposes into separable functions,

$$(6) \quad \phi(\omega) = \sum_{s \in \mathcal{S}} p_s \phi_s(\omega_s),$$

where for each $s \in \mathcal{S}$,

$$(7) \quad \phi_s(\omega_s) := \min_{x,y} \{(c + \omega_s)^\top x + q_s^\top y : (x, y) \in K_s\}.$$

The reformulation (6) is the basis for parallelizable approaches for computing dual bounds that are used, for example, in the dual decomposition methods developed in [9, 23].

For any choice of $\omega = (\omega_1, \dots, \omega_{|\mathcal{S}|})$, it is well known that the value of the Lagrangian dual function provides a lower bound on the optimal solution to (1): $\phi(\omega) \leq \zeta^{SMIP}$. The problem of finding the best such lower bound is the *Lagrangian dual problem*:

$$(8) \quad \zeta^{LD} := \sup_{\omega} \left\{ \phi(\omega) : \sum_{s \in \mathcal{S}} p_s \omega_s = 0 \right\}.$$

The primary contribution of this work is a new and effective method for solving (8), thus enabling a practical and efficient computation of high-quality lower bounds for ζ^{SMIP} .

The function $\phi(\omega)$ is a piecewise-affine concave function, and many methods are known for maximizing such functions. These methods include the subgradient method [35], the augmented Lagrangian (AL) method [16, 31], and the alternating direction method of multipliers (ADMM) [14, 12, 8]. The subgradient method has mainly theoretical significance, since it is difficult to develop reliable and efficient step-size rules for the dual variables ω (see, e.g., section 7.1.1 of [33]). As iterative primal-dual approaches, methods based on the AL method or ADMM are more effective in practice. However, in the context of SMIP, both methods require convexification of the constraints K_s , $s \in \mathcal{S}$, to have a meaningful theoretical support for convergence to the best lower bound value ζ^{LD} . Furthermore, both methods require the solution of additional mixed-integer linear programming (MILP) subproblems in order to recover the Lagrangian lower bounds associated with the dual values, ω [15]. ADMM has a more straightforward potential for decomposability and parallelization than the AL method, and so in this work we develop a theoretically supported modification of a method based on ADMM.

When specialized to the deterministic equivalent problem (2) in the context of stochastic programming, ADMM is referred to as progressive hedging (PH) [32, 39]. When the sets K_s , $s \in \mathcal{S}$, are convex, the limit points of the sequence of solution-multiplier pairs $\{(x^k, y^k, z^k), \omega^k\}_{k=1}^{\infty}$ generated by PH are saddle points of the deterministic equivalent problem (2), whenever such saddle points exist. When the constraints $(x_s, y_s) \in K_s$, $s \in \mathcal{S}$, enforce nontrivial mixed-integer restrictions, the set K_s is not convex and PH becomes a heuristic approach with no guarantees of convergence [21]. Nevertheless, some measure of success in practice has been observed in [39] while applying PH to problems of the form (3). More recently, [15] showed that valid Lagrangian lower bounds can be calculated from the iterates of the PH algorithm when the sets K_s are not convex. However, their implementation of the

algorithm does not offer any guarantee that the lower bounds will converge to the optimal value ζ^{LD} . Moreover, additional computational effort, in solving additional MILP subproblems, must be expended in order to compute the lower bound. Our contribution is to extend the PH-based approach in [15], creating an algorithm whose lower bound values converge to ζ^{LD} in theory and for which lower bound calculations do not require additional computational effort. Computational results in section 4 demonstrate that the new method outperforms the existing PH-based method, in terms of both quality of bound and efficiency of computation.

To motivate our approach, we first consider the application of PH to the following well-known primal characterization of ζ^{LD} :

$$(9) \quad \zeta^{LD} = \min_{x,y,z} \left\{ \sum_{s \in \mathcal{S}} p_s (c^\top x_s + q_s^\top y_s) : (x_s, y_s) \in \text{conv}(K_s), x_s = z \ \forall s \in \mathcal{S} \right\},$$

where $\text{conv}(K_s)$ denotes the convex hull of K_s for each $s \in \mathcal{S}$. (See, for example, Theorem 6.2 of [25].)

The sequence of Lagrangian bounds $\{\phi(\omega^k)\}$ generated by the application of PH to (9) is known to be convergent. Thus, the value of the Lagrangian dual problem (ζ^{LD}) may, in theory, be computed by applying PH to (9). However, in practice, an explicit polyhedral description of $\text{conv}(K_s)$, $s \in \mathcal{S}$, is generally not available, thus raising the issue of implementability.

The absence of such an explicit description motivates an application of a solution approach to the PH primal update step that iteratively constructs an improved inner approximation of each $\text{conv}(K_s)$, $s \in \mathcal{S}$. For this purpose, we apply a solution approach to the PH primal update problem that is based on the Frank–Wolfe (FW) method [13]. Our approach has the additional benefit of providing Lagrangian bounds at no additional computational cost.

One simple, theoretically supported integration of an FW-like method and PH is realized by having the PH primal updates computed using a method called the simplicial decomposition method (SDM) [17, 38]. SDM is an extension of the FW method that makes use of progressively improving inner approximations to each set $\text{conv}(K_s)$, $s \in \mathcal{S}$. The finite optimal convergence of each application of SDM follows directly from the polyhedral structure $\text{conv}(K_s)$ and the (practically reasonable) assumption that $\text{conv}(K_s)$ is bounded for each $s \in \mathcal{S}$.

For computing improvements in the Lagrangian bound efficiently, convergence of SDM to the optimal solution of the subproblem is too costly and not necessary. We thus develop a modified integration whose theoretically supported convergence analysis is based not on the optimal convergence of SDM, but rather on its ability to adequately extend the inner approximations of each $\text{conv}(K_s)$, $s \in \mathcal{S}$.

The main contribution of this paper is the development, convergence analysis, and application of a new algorithm, called FW-PH, which is used to compute high-quality Lagrangian bounds for SMIPs efficiently and with a high potential for parallelization. FW-PH is efficient in that, under mild assumptions, each dual update and Lagrangian bound computation may be obtained by solving, for each $s \in \mathcal{S}$, just one MILP problem and one continuous convex quadratic problem. In contrast, each dual update of PH requires the solution of a mixed-integer quadratic programming (MIQP) subproblem for each $s \in \mathcal{S}$, and each PH Lagrangian bound computation requires the solution of one MILP subproblem for each $s \in \mathcal{S}$. In our convergence analysis, conditions are provided under which the sequence of Lagrangian bounds generated by FW-PH converges to the optimal Lagrangian bound ζ^{LD} . To the best of our knowledge, the

combination of PH and FW in a manner that is theoretically supported, computationally efficient, and parallelizable is new, in spite of the convergence analyses of both PH and FW being well developed. We also perform an experimental assessment of alternative heuristic strategies that can be employed in a straightforward manner to recover feasible solutions for problem (3) from the solution obtained by FW-PH for problem (9).

This paper is organized as follows. In section 2, we present the theoretical background of PH and a brief technical lemma regarding the inner approximations generated by SDM; this background is foundational to the proposed FW-PH method. In section 3, we present the FW-PH method and a convergence analysis. We also present in this section heuristic strategies that can be employed to generate primal feasible first-stage solutions. The results of numerical experiments comparing the Lagrangian bounds computed with PH and those with FW-PH are presented in section 4. We also provide a comparison of the primal solutions obtained employing the heuristics described. We conclude in section 5 with a discussion of the results obtained and with suggested directions for further research.

2. Progressive hedging and Frank–Wolfe-based methods. The *augmented Lagrangian (AL) dual function* based on the relaxation of the nonanticipativity constraints $x_s = z$, $s \in \mathcal{S}$, is

$$L^\rho(x, y, z, \omega) := \sum_{s \in \mathcal{S}} p_s L_s^\rho(x_s, y_s, z, \omega_s),$$

where

$$L_s^\rho(x_s, y_s, z, \omega_s) := c^\top x_s + q_s^\top y_s + \omega_s^\top (x_s - z) + \frac{\rho}{2} \|x_s - z\|_2^2$$

and $\rho > 0$ is a penalty parameter. By changing the feasible region, denoted here by D_s , $s \in \mathcal{S}$, the AL dual problem (i.e., the Lagrangian dual problem in which the Lagrangian dual function is replaced by its augmented version) can be used in a progressive hedging (PH) approach to solve either problem (3) or problem (9). Pseudocode for the PH algorithm is given in Algorithm 1.

In Algorithm 1, $k_{max} > 0$ is the maximum number of iterations and $\epsilon > 0$ parameterizes the convergence tolerance. The initialization of lines 3–8 provides an initial target primal value z^0 and dual values ω_s^1 , $s \in \mathcal{S}$, for the main iterations $k \geq 1$. Also, an initial Lagrangian bound ϕ^0 can be computed from this initialization.

For $\epsilon > 0$, the Algorithm 1 termination criterion $\sqrt{\sum_{s \in \mathcal{S}} p_s \|x_s^k - z^{k-1}\|_2^2} < \epsilon$ is motivated by the addition of the squared norms of the primal and dual residuals associated with problem (9). These residuals, as developed in section 3.3 of [8] within the more general context of ADMM, are measures of how close (x^k, y^k, z^k) come to satisfying the necessary and sufficient conditions of optimality for problem (9). Hence, we enforce an adequate vanishing of primal and dual feasibility residuals, which ultimately implies the vanishing of primal (and dual) objective value suboptimality [8]. In summing the squared norm primal residuals $p_s \|x_s^k - z^k\|_2^2$, $s \in \mathcal{S}$, and the squared norm dual residual $\|z^k - z^{k-1}\|_2^2$, we have

$$(10) \quad \sum_{s \in \mathcal{S}} p_s \left[\|x_s^k - z^k\|_2^2 + \|z^k - z^{k-1}\|_2^2 \right] = \sum_{s \in \mathcal{S}} p_s \|x_s^k - z^{k-1}\|_2^2.$$

The equality in (10) follows since, for each $s \in \mathcal{S}$, the cross term resulting from the expansion of the squared norm $\|(x_s^k - z^k) + (z^k - z^{k-1})\|_2^2$ vanishes; this is seen in the equality $\sum_{s \in \mathcal{S}} p_s (x_s^k - z^k) = 0$ due to the construction of z^k .

Algorithm 1 PH applied to problem (3) ($D_s = K_s$) or (9) ($D_s = \text{conv}(K_s)$).

```

1: Precondition:  $\sum_{s \in \mathcal{S}} p_s \omega_s^0 = 0$ 
2: function PH( $\omega^0, \rho, k_{max}, \epsilon$ )
3:   for  $s \in \mathcal{S}$  do
4:      $(x_s^0, y_s^0) \in \text{argmin}_{x,y} \{(c + \omega_s^0)^\top x + q_s^\top y : (x, y) \in D_s\}$ 
5:   end for
6:    $\phi^0 \leftarrow \sum_{s \in \mathcal{S}} p_s [(c + \omega_s^0)^\top x_s^0 + q_s^\top y_s^0]$ 
7:    $z^0 \leftarrow \sum_{s \in \mathcal{S}} p_s x_s^0$ 
8:    $\omega_s^1 \leftarrow \omega_s^0 + \rho(x_s^0 - z^0)$  for all  $s \in \mathcal{S}$ 
9:   for  $k = 1, \dots, k_{max}$  do
10:    for  $s \in \mathcal{S}$  do
11:       $\phi_s^k \leftarrow \min_{x,y} \{(c + \omega_s^k)^\top x + q_s^\top y : (x, y) \in D_s\}$ 
12:       $(x_s^k, y_s^k) \in \text{argmin}_{x,y} \{L_s^\rho(x, y, z^{k-1}, \omega_s^k) : (x, y) \in D_s\}$ 
13:    end for
14:     $\phi^k \leftarrow \sum_{s \in \mathcal{S}} p_s \phi_s^k$ 
15:     $z^k \leftarrow \sum_{s \in \mathcal{S}} p_s x_s^k$ 
16:    if  $\sqrt{\sum_{s \in \mathcal{S}} p_s \|x_s^k - z^{k-1}\|_2^2} < \epsilon$  then
17:      return  $(x^k, y^k, z^k, \omega^k, \phi^k)$ 
18:    end if
19:     $\omega_s^{k+1} \leftarrow \omega_s^k + \rho(x_s^k - z^k)$  for all  $s \in \mathcal{S}$ 
20:  end for
21:  return  $(x^{k_{max}}, y^{k_{max}}, z^{k_{max}}, \omega^{k_{max}}, \phi^{k_{max}})$ 
22: end function

```

The line 11 subproblem of Algorithm 1 is an addition to the original PH algorithm. Its purpose is to compute Lagrangian bounds (line 14) from the current dual solution ω^k [15]. Thus, the bulk of computational effort in Algorithm 1 applied to problem (3) (the case with $D_s = K_s$) resides in computing solutions to the MILP (line 11) and MIQP (line 12) subproblems. Note that line 11 (and line 14) may be omitted if the corresponding Lagrangian bound for ω^k is not desired.

2.1. Convergence of PH. The following proposition addresses the convergence of PH applied to problem (9).

PROPOSITION 2.1. *Assume that problem (9) is feasible with $\text{conv}(K_s)$ bounded for each $s \in \mathcal{S}$, and let Algorithm 1 be applied to problem (9) (so that $D_s = \text{conv}(K_s)$ for each $s \in \mathcal{S}$) with tolerance $\epsilon = 0$ for each $k \geq 1$. Then, the limit $\lim_{k \rightarrow \infty} \omega^k = \omega^*$ exists and, furthermore,*

1. $\lim_{k \rightarrow \infty} \sum_{s \in \mathcal{S}} p_s (c^\top x_s^k + q_s^\top y_s^k) = \zeta^{LD}$,
2. $\lim_{k \rightarrow \infty} \phi(\omega^k) = \zeta^{LD}$,
3. $\lim_{k \rightarrow \infty} (x_s^k - z^k) = 0$ for each $s \in \mathcal{S}$,

and each limit point $((x_s^*, y_s^*)_{s \in \mathcal{S}}, z^*)$ is an optimal solution for (9).

Proof. Since the constraint sets $D_s = \text{conv}(K_s)$, $s \in \mathcal{S}$, are bounded, and problem (9) is feasible, problem (9) has an optimal solution $((x_s^*, y_s^*)_{s \in \mathcal{S}}, z^*)$ with optimal value ζ^{LD} . The feasibility of problem (9), the linearity of its objective function, and the bounded polyhedral structure of its constraint set $D_s = \text{conv}(K_s)$, $s \in \mathcal{S}$, imply that the hypotheses for PH convergence to the optimal solution are met (See Theorem 5.1 of [32]). Therefore, $\{\omega^k\}$ converges to some ω^* , $\lim_{k \rightarrow \infty} \sum_{s \in \mathcal{S}} p_s (c^\top x_s^k + q_s^\top y_s^k) =$

ζ^{LD} , $\lim_{k \rightarrow \infty} \phi(\omega^k) = \zeta^{LD}$, and $\lim_{k \rightarrow \infty} (x_s^k - z^k) = 0$ for each $s \in \mathcal{S}$ all hold. The boundedness of each $D_s = \text{conv}(K_s)$, $s \in \mathcal{S}$, furthermore implies the existence of limit points $((x_s^*, y_s^*)_{s \in \mathcal{S}}, z^*)$ of $\{((x_s^k, y_s^k)_{s \in \mathcal{S}}, z^k)\}$, which are optimal solutions for (9). \square

Note that the convergence in Proposition 2.1 applies to the continuous problem (9) but *not* to the mixed-integer problem (3). In problem (3), the constraint sets K_s , $s \in \mathcal{S}$, are not convex, so there is no guarantee that Algorithm 1 will converge when applied to (3). However, the application of PH to problem (9) requires, in line 12, the optimization of the AL over the sets $\text{conv}(K_s)$, $s \in \mathcal{S}$, for which an explicit linear description is unlikely to be known. In the next section, we demonstrate how to circumvent this difficulty by constructing inner approximations of the polyhedral sets $\text{conv}(K_s)$, $s \in \mathcal{S}$.

2.2. A Frank–Wolfe approach based on simplicial decomposition. To use Algorithm 1 to solve (9) requires a method for solving the subproblem

$$(11) \quad (x_s^k, y_s^k) \in \underset{x, y}{\text{argmin}} \{L_s^\rho(x, y, z^{k-1}, \omega_s^k) : (x, y) \in \text{conv}(K_s)\}$$

appearing in line 12 of the algorithm. Although an explicit description of $\text{conv}(K_s)$ is not readily available, if we have a *linear* objective function, then we can replace $\text{conv}(K_s)$ with K_s (solving one MIP problem per scenario). This motivates the application of an FW algorithm for solving (11), since the FW algorithm solves a sequence of problems in which the nonlinear objective is linearized using a first-order approximation.

The simplicial decomposition method (SDM) is an extension of the FW method, where the line searches of FW are replaced by searches over polyhedral inner approximations. SDM can be applied to solve a feasible, bounded problem of the general form

$$(12) \quad \zeta^{FW} := \min_x \{f(x) : x \in D\},$$

with nonempty compact convex set D and continuously differentiable convex function f . Generically, given a current solution x^{t-1} and inner approximation $D^{t-1} \subseteq D$, iteration t of the SDM consists of solving

$$\hat{x} \in \underset{x}{\text{argmin}} \{\nabla_x f(x^{t-1})^\top x : x \in D\},$$

updating the inner approximation as $D^t \leftarrow \text{conv}(D^{t-1} \cup \{\hat{x}\})$, and finally choosing

$$x^t \in \underset{x}{\text{argmin}} \{f(x) : x \in D^t\}.$$

The algorithm terminates when the bound gap is small, specifically, when

$$\Gamma^t := -\nabla_x f(x^{t-1})^\top (\hat{x} - x^{t-1}) \leq \tau,$$

where $\tau \geq 0$ is a given tolerance.

The application of SDM to solve problem (11), i.e., to minimize $L_s^\rho(x, y, z, \omega_s)$ over $(x, y) \in \text{conv}(K_s)$ for a given $s \in \mathcal{S}$, is presented in Algorithm 2. Here, t_{max} is the maximum number of iterations and $\tau > 0$ is a convergence tolerance. Γ^t is the bound gap used to measure closeness to optimality, and ϕ_s is used to compute a Lagrangian bound as described in the next section. The inner approximation to

Algorithm 2 SDM applied to problem (11).

```

1: Precondition:  $V_s^0 \subset \text{conv}(K_s)$  and  $z = \sum_{s \in S} p_s x_s^0$ 
2: function SDM( $V_s^0, x_s^0, \omega_s, z, t_{max}, \tau$ )
3:   for  $t = 1, \dots, t_{max}$  do
4:      $\hat{\omega}_s^t \leftarrow \omega_s + \rho(x_s^{t-1} - z)$ 
5:      $(\hat{x}_s, \hat{y}_s) \in \text{argmin}_{x,y} \{ (c + \hat{\omega}_s^t)^\top x + q_s^\top y : (x, y) \in \mathcal{V}(\text{conv}(K_s)) \}$ 
6:     if  $t = 1$  then
7:        $\phi_s \leftarrow (c + \hat{\omega}_s^t)^\top \hat{x}_s + q_s^\top \hat{y}_s$ 
8:     end if
9:      $\Gamma^t \leftarrow -[(c + \hat{\omega}_s^t)^\top (\hat{x}_s - x_s^{t-1}) + q_s^\top (\hat{y}_s - y_s^{t-1})]$ 
10:     $V_s^t \leftarrow V_s^{t-1} \cup \{(\hat{x}_s, \hat{y}_s)\}$ 
11:     $(x_s^t, y_s^t) \in \text{argmin}_{x,y} \{ L_s^\rho(x, y, z, \omega_s) : (x, y) \in \text{conv}(V_s^t) \}$ 
12:    if  $\Gamma^t \leq \tau$  then
13:      return  $(x_s^t, y_s^t, V_s^t, \phi_s)$ 
14:    end if
15:  end for
16:  return  $(x_s^{t_{max}}, y_s^{t_{max}}, V_s^{t_{max}}, \phi_s)$ 
17: end function

```

$\text{conv}(K_s)$ at iteration $t \geq 1$ takes the form $\text{conv}(V_s^t)$, where V_s^t is a finite set of points with $V_s^t \subset \text{conv}(K_s)$. The points added by Algorithm 2 to the initial set, V_s^0 , to form V_s^t are all in K_s : here $\mathcal{V}(\text{conv}(K_s))$ is the set of extreme points of $\text{conv}(K_s)$ and, of course, $\mathcal{V}(\text{conv}(K_s)) \subseteq K_s$.

Observe that

$$\nabla_{(x,y)} L_s^\rho(x, y, z, \omega_s) |_{(x,y)=(x_s^{t-1}, y_s^{t-1})} = \begin{bmatrix} c + \omega_s + \rho(x_s^{t-1} - z) \\ q_s \end{bmatrix} = \begin{bmatrix} c + \hat{\omega}_s \\ q_s \end{bmatrix},$$

with $\hat{\omega}_s = \omega_s + \rho(x_s^{t-1} - z)$, and so the optimization at line 5 is minimizing the gradient approximation to $L_s^\rho(x, y, z, \omega_s)$ at the point (x_s^{t-1}, y_s^{t-1}) . Since this is a linear objective function, optimization over $\mathcal{V}(\text{conv}(K_s))$ can be accomplished by optimization over K_s (see, e.g., section I.4, Theorem 6.3 of [25]). Hence line 5 requires a solution of a single-scenario MILP.

The optimization at line 11 can be accomplished by expressing (x, y) as a convex combination of the finite set of points, V_s^t , where the weights $a \in \mathbb{R}^{|V_s^t|}$ in the convex combination are now also decision variables. That is, the line 11 problem is solved with a solution to the following convex continuous quadratic subproblem:

$$(13) \quad (x_s^t, y_s^t, a) \in \text{argmin}_{x,y,a} \left\{ \begin{array}{l} L_s^\rho(x, y, z, \omega_s) : (x, y) = \sum_{(\hat{x}^i, \hat{y}^i) \in V_s^t} a_i (\hat{x}^i, \hat{y}^i), \\ \sum_{i=1, \dots, |V_s^t|} a_i = 1, \text{ and } a_i \geq 0 \text{ for } i = 1, \dots, |V_s^t| \end{array} \right\}.$$

For implementational purposes, the x and y variables may be substituted out of the objective of problem (13), leaving a as the only decision variable, with the only constraints being nonnegativity of the a components and the requirement that they sum to 1.

SDM is known to terminate finitely with an optimal solution when D is polyhedral [17], so the primal update step line 12, Algorithm 1 with $D_s = \text{conv}(K_s)$ could be accomplished with SDM, resulting in an algorithm that converges to a solution to problem (9). This solution, despite not being feasible for problem (3) in general (as it

typically does not observe integrality requirements), gives the Lagrangian dual bound ζ^{LD} . However, since each inner iteration of line 5, Algorithm 2 requires the solution of a MILP, using t_{\max} large enough to ensure SDM terminates optimally is not efficient for our purpose of computing Lagrangian bounds. In the next section, we give an adaptation of the algorithm that requires the solution of only *one* MILP subproblem per scenario at each major iteration of the PH algorithm.

3. The FW-PH method. In order to make SDM efficient when used with PH to solve problem (9), the minimization of the augmented Lagrangian dual problem can be solved approximately. This insight can greatly reduce the number of MILP subproblems solved at each inner iteration and forms the basis of our algorithm FW-PH. Convergence of FW-PH relies on the following lemma, which states an important expansion property of the inner approximations employed by SDM.

LEMMA 3.1. *For any scenario $s \in \mathcal{S}$ and iteration $k \geq 1$, let Algorithm 2 be applied to the minimization problem (11) for any $t_{\max} \geq 2$. For $1 \leq t < t_{\max}$, if*

$$(14) \quad (x_s^t, y_s^t) \notin \underset{x, y}{\operatorname{argmin}} \{L_\rho^s(x, y, z^{k-1}, \omega_s^k) : (x, y) \in \operatorname{conv}(K_s)\}$$

holds, then $\operatorname{conv}(V_s^{t+1}) \supset \operatorname{conv}(V_s^t)$.

Proof. For $s \in \mathcal{S}$ and $k \geq 1$ fixed, we know that by construction

$$(x_s^t, y_s^t) \in \underset{x, y}{\operatorname{argmin}} \{L_\rho^s(x, y, z^{k-1}, \omega_s^k) : (x, y) \in \operatorname{conv}(V_s^t)\}$$

for $t \geq 1$. Given the convexity of $(x, y) \mapsto L_\rho^s(x, y, z^{k-1}, \omega_s^k)$ and the convexity of $\operatorname{conv}(V_s^t)$, the necessary and sufficient condition for optimality

$$(15) \quad \nabla_{(x, y)} L_\rho^s(x_s^t, y_s^t, z^{k-1}, \omega_s^k) \begin{bmatrix} x - x_s^t \\ y - y_s^t \end{bmatrix} \geq 0 \quad \text{for all } (x, y) \in \operatorname{conv}(V_s^t)$$

is satisfied. By assumption, condition (14) is satisfied, $\operatorname{conv}(K_s)$ is likewise convex, and so the resulting *nonsatisfaction* of the necessary and sufficient condition of optimality for the problem in (14) takes the form

$$(16) \quad \min_{x, y} \left\{ \nabla_{(x, y)} L_\rho^s(x_s^t, y_s^t, z^{k-1}, \omega_s^k) \begin{bmatrix} x - x_s^t \\ y - y_s^t \end{bmatrix} : (x, y) \in \operatorname{conv}(K_s) \right\} < 0.$$

In fact, during SDM iteration $t + 1$, an optimal solution (\hat{x}_s, \hat{y}_s) to the problem in condition (16) is computed in line 5 of Algorithm 2. Therefore, by the satisfaction of condition (15) and the optimality of (\hat{x}_s, \hat{y}_s) for the problem of condition (16), which is also satisfied, we have $(\hat{x}_s, \hat{y}_s) \notin \operatorname{conv}(V_s^t)$. By construction, $V_s^{t+1} \leftarrow V_s^t \cup \{(\hat{x}_s, \hat{y}_s)\}$, so that $\operatorname{conv}(V_s^{t+1}) \supset \operatorname{conv}(V_s^t)$ must hold. \square

3.1. Convergence of FW-PH. The FW-PH algorithm is stated in pseudocode form in Algorithm 3. Similar to Algorithm 1, the parameter ϵ is a convergence tolerance, and k_{\max} is the maximum number of (outer) iterations. The parameter t_{\max} is the maximum number of (inner) SDM iterations in Algorithm 2.

The parameter $\alpha \in \mathbb{R}$ affects the initial linearization point \tilde{x}_s of the SDM method. Any value $\alpha \in \mathbb{R}$ may be used, but the use of $\tilde{x}_s = (1 - \alpha)z^{k-1} + \alpha x_s^{k-1}$ in line 6 is a crucial component in the efficiency of the FW-PH algorithm, as it enables the computation of a valid dual bound, ϕ^k , at *each* iteration of FW-PH without the need for additional MILP subproblem solutions. Specifically, we have the following result.

Algorithm 3 FW-PH applied to problem (9).

```

1: function FW-PH( $(V_s^0)_{s \in \mathcal{S}}, (x_s^0, y_s^0)_{s \in \mathcal{S}}, \omega^0, \rho, \alpha, \epsilon, k_{max}, t_{max}$ )
2:    $z^0 \leftarrow \sum_{s \in \mathcal{S}} p_s x_s^0$ 
3:    $\omega_s^1 \leftarrow \omega_s^0 + \rho(x_s^0 - z^0)$  for  $s \in \mathcal{S}$ 
4:   for  $k = 1, \dots, k_{max}$  do
5:     for  $s \in \mathcal{S}$  do
6:        $\tilde{x}_s \leftarrow (1 - \alpha)z^{k-1} + \alpha x_s^{k-1}$ 
7:        $[x_s^k, y_s^k, V_s^k, \phi_s^k] \leftarrow \text{SDM}(V_s^{k-1}, \tilde{x}_s, \omega_s^k, z^{k-1}, t_{max}, 0)$ 
8:     end for
9:      $\phi^k \leftarrow \sum_{s \in \mathcal{S}} p_s \phi_s^k$ 
10:     $z^k \leftarrow \sum_{s \in \mathcal{S}} p_s x_s^k$ 
11:    if  $\sqrt{\sum_{s \in \mathcal{S}} p_s \|x_s^k - z^{k-1}\|_2^2} < \epsilon$  then
12:      return  $((x_s^k, y_s^k)_{s \in \mathcal{S}}, z^k, \omega^k, \phi^k)$ 
13:    end if
14:     $\omega_s^{k+1} \leftarrow \omega_s^k + \rho(x_s^k - z^k)$  for  $s \in \mathcal{S}$ 
15:  end for
16:  return  $((x_s^{k_{max}}, y_s^{k_{max}})_{s \in \mathcal{S}}, z^{k_{max}}, \omega^{k_{max}}, \phi^{k_{max}})$ 
17: end function

```

PROPOSITION 3.2. Assume that the precondition $\sum_{s \in \mathcal{S}} p_s \omega_s^0 = 0$ holds for Algorithm 3. At each iteration $k \geq 1$ of Algorithm 3, the value, ϕ^k , calculated at line 9, is the value of the Lagrangian dual function $\phi(\cdot)$ evaluated at a Lagrangian dual feasible point, and hence provides a finite lower bound on ζ^{LD} .

Proof. Since $\sum_{s \in \mathcal{S}} p_s \omega_s^0 = 0$ holds and, by construction, $0 = \sum_{s \in \mathcal{S}} p_s (x_s^0 - z^0)$, we have $\sum_{s \in \mathcal{S}} p_s \omega_s^1 = 0$ also. We proceed by induction on $k \geq 1$. At iteration k , the problem solved for each $s \in \mathcal{S}$ at line 5 in the first iteration ($t = 1$) of Algorithm 2 may be solved with the same optimal value by exchanging $\mathcal{V}(\text{conv}(K_s))$ for K_s ; this follows from the linearity of the objective function. Thus, an optimal solution computed at line 5 may be used in the computation of $\phi_s(\tilde{\omega}_s^k)$ carried out in line 7, where

$$\begin{aligned} \tilde{\omega}_s^k &:= \hat{\omega}_s^1 = \omega_s^k + \rho(\tilde{x}_s - z^{k-1}) = \omega_s^k + \rho((1 - \alpha)z^{k-1} + \alpha x_s^{k-1} - z^{k-1}) \\ &= \omega_s^k + \alpha \rho(x_s^{k-1} - z^{k-1}). \end{aligned}$$

By construction, we have at each iteration $k \geq 1$ in Algorithm 3 that

$$\sum_{s \in \mathcal{S}} p_s (x_s^{k-1} - z^{k-1}) = 0 \quad \text{and} \quad \sum_{s \in \mathcal{S}} p_s \omega_s^k = 0,$$

which establishes that $\sum_{s \in \mathcal{S}} p_s \tilde{\omega}_s^k = 0$. Thus, $\tilde{\omega}^k$ is feasible for the Lagrangian dual problem, so that $\phi(\tilde{\omega}^k) = \sum_{s \in \mathcal{S}} p_s \phi_s^k$, and, since each ϕ_s^k is the optimal value of a bounded and feasible mixed-integer linear program, we have $-\infty < \phi(\tilde{\omega}^k) < \infty$. \square

We establish convergence of Algorithm 3 for any $\alpha \in \mathbb{R}$ and $t_{max} \geq 1$. For the special case in which we perform only one iteration of SDM for each outer iteration ($t_{max} = 1$), we require the additional assumption that the initial scenario vertex sets share a common point. More precisely, we require the assumption

$$(17) \quad \bigcap_{s \in \mathcal{S}} \text{Proj}_x(\text{conv}(V_s^0)) \neq \emptyset,$$

which can, in practice, be effectively handled through appropriate initialization, under the standard assumption of relatively complete recourse: for all $x \in X$ and $s \in \mathcal{S}$ there exists y_s such that $(x, y_s) \in K_s$. We describe one such initialization approach in section 4.

PROPOSITION 3.3. *Let the convexified separable deterministic equivalent SMIP (9) have an optimal solution, and let Algorithm 3 be applied to (9) with $k_{max} = \infty$, $\epsilon = 0$, $\alpha \in \mathbb{R}$, and $t_{max} \geq 1$. If either $t_{max} \geq 2$ or (17) holds, then $\lim_{k \rightarrow \infty} \phi^k = \zeta^{LD}$.*

Proof. First note that for any $t_{max} \geq 1$ the sequence of inner approximations $\text{conv}(V_s^k)$, $s \in \mathcal{S}$, will stabilize in that, for some threshold $0 \leq \bar{k}_s$, we have, for all $k \geq \bar{k}_s$,

$$(18) \quad \text{conv}(V_s^k) =: \bar{D}_s \subseteq \text{conv}(K_s).$$

This follows due to the assumption that each expansion of the inner approximations $\text{conv}(V_s^k)$ takes the form $V_s^k \leftarrow V_s^{k-1} \cup \{(\hat{x}_s, \hat{y}_s)\}$, where (\hat{x}_s, \hat{y}_s) is a vertex of $\text{conv}(K_s)$. Since each polyhedron $\text{conv}(K_s)$, $s \in \mathcal{S}$, has only a finite number of such vertices, the stabilization (18) must occur at some $\bar{k}_s < \infty$.

For $t_{max} \geq 2$, the stabilizations (18), $s \in \mathcal{S}$, are reached at some iteration $\bar{k} := \max_{s \in \mathcal{S}} \{\bar{k}_s\}$. Noting that $\bar{D}_s = \text{conv}(V_s^k)$ for $k > \bar{k}$ we must have

$$(19) \quad (x_s^k, y_s^k) \in \underset{x, y}{\text{argmin}} \{L_s^\rho(x, y, z^{k-1}, \omega_s^k) : (x, y) \in \text{conv}(K_s)\}.$$

Otherwise, due to Lemma 3.1, the call to SDM on line 7 must return $V_s^k \supset V_s^{k-1}$, contradicting the finite stabilization (18). Therefore, the $k \geq \bar{k}$ iterations of Algorithm 3 are identical to Algorithm 1 iterations, and so Proposition 2.1 implies that $\lim_{k \rightarrow \infty} x_s^k - z^k = 0$, $s \in \mathcal{S}$, and $\lim_{k \rightarrow \infty} \phi(\omega^k) = \zeta^{LD}$. By the continuity of $\omega \mapsto \phi_s(\omega)$ for each $s \in \mathcal{S}$, we have $\lim_{k \rightarrow \infty} \phi^k = \lim_{k \rightarrow \infty} \sum_{s \in \mathcal{S}} p_s \phi_s(\omega_s^k + \alpha(x_s^{k-1} - z^{k-1})) = \lim_{k \rightarrow \infty} \sum_{s \in \mathcal{S}} p_s \phi_s(\omega_s^k) = \lim_{k \rightarrow \infty} \phi(\omega^k) = \zeta^{LD}$ for all $\alpha \in \mathbb{R}$.

In the $t_{max} = 1$ case, we have at each iteration $k \geq 1$ the optimality

$$(x_s^k, y_s^k) \in \underset{x, y}{\text{argmin}} \{L_s^\rho(x_s, y_s, z^{k-1}, \omega_s^k) : (x_s, y_s) \in \text{conv}(V_s^k)\}.$$

By the definition of stabilization (18), the iterations $k \geq \bar{k}$ of Algorithm 3 are identical to PH iterations applied to the restricted problem

$$(20) \quad \min_{x, y, z} \left\{ \sum_{s \in \mathcal{S}} p_s (c^\top x_s + q_s^\top y_s) : (x_s, y_s) \in \bar{D}_s \ \forall s \in \mathcal{S}, \ x_s = z \ \forall s \in \mathcal{S} \right\}.$$

We have initialized the sets $(V_s^0)_{s \in \mathcal{S}}$ such that $\bigcap_{s \in \mathcal{S}} \text{Proj}_x \text{conv}(V_s^0) \neq \emptyset$, so since the inner approximations to $\text{conv}(K_s)$ only expand in the algorithm, we have that $\bigcap_{s \in \mathcal{S}} \text{Proj}_x(\bar{D}_s) \neq \emptyset$. Therefore, problem (20) is a feasible and bounded linear program, and so the PH convergence described in Proposition 2.1 with $D_s = \bar{D}_s$, $s \in \mathcal{S}$ holds for its application to problem (20). That is, for each $s \in \mathcal{S}$, we have (1) $\lim_{k \rightarrow \infty} \omega_s^k = \omega_s^*$ and $\lim_{k \rightarrow \infty} (x_s^k - z^k) = 0$; and (2) for all limit points $((x_s^*, y_s^*)_{s \in \mathcal{S}}, z^*)$, we have the feasibility and optimality of the limit points, which implies $x_s^* = z^*$ and

$$(21) \quad \min_{x, y} \{(c + \omega_s^*)^\top (x - x^*) + q_s^\top (y - y^*) : (x, y) \in \bar{D}_s\} = 0.$$

Next, for each $s \in \mathcal{S}$ the compactness of $\text{conv}(K_s) \supseteq \overline{D}_s$, the continuity of the minimum value function

$$\omega \mapsto \min_{x,y} \{(c + \omega)^\top x + q_s^\top y : (x, y) \in \overline{D}_s\}$$

over $\omega \in \{\omega : \sum_{s \in \mathcal{S}} p_s \omega_s = 0\}$, and the limit $\lim_{k \rightarrow \infty} \tilde{\omega}_s^{k+1} = \lim_{k \rightarrow \infty} \omega_s^{k+1} + \alpha \rho(x_s^k - z^k) = \omega_s^*$, together imply that

$$(22) \quad \lim_{k \rightarrow \infty} \min_{x,y} \{(c + \tilde{\omega}_s^{k+1})^\top (x - x^k) + q_s^\top (y - y^k) : (x, y) \in \overline{D}_s\} = 0.$$

Recall that $\tilde{\omega}_s^k = \omega_s^k + \rho \alpha (x_s^{k-1} - z^{k-1})$ is the $t = 1$ value of $\hat{\omega}_s^t$ defined in line 4 of Algorithm 2. Thus, for $k + 1 > \bar{k}$, we have due to the stabilization (18) that

$$(23) \quad \min_{x,y} \{(c + \tilde{\omega}_s^{k+1})^\top (x - x^k) + q_s^\top (y - y^k) : (x, y) \in \overline{D}_s\} = \min_{x,y} \{(c + \tilde{\omega}_s^{k+1})^\top (x - x^k) + q_s^\top (y - y^k) : (x, y) \in \text{conv}(K_s)\}$$

If (23) does not hold, then the inner approximation expansion $\overline{D}_s \subset \text{conv}(V_s^{k+1})$ must occur, since a point $(\hat{x}_s, \hat{y}_s) \in \text{conv}(K_s)$ that can be strictly separated from \overline{D}_s would have been discovered during the iteration $k + 1$ execution of Algorithm 2, line 5, $t = 1$. The expansion $\overline{D}_s \subset \text{conv}(V_s^{k+1})$ contradicts the finite stabilization (18), and so (23) holds. Therefore, the equalities (22) and (23) imply that

$$(24) \quad \lim_{k \rightarrow \infty} \min_{x,y} \{(c + \tilde{\omega}_s^{k+1})^\top (x - x^k) + q_s^\top (y - y^k) : (x, y) \in \text{conv}(K_s)\} = 0.$$

Our argument has shown that for all limit points (x_s^*, y_s^*) , $s \in \mathcal{S}$, the stationarity condition

$$(25) \quad (c + \omega_s^*)^\top (x - x_s^*) + q_s^\top (y - y_s^*) \geq 0 \quad \forall (x, y) \in \text{conv}(K_s)$$

is satisfied, which together with the feasibility $x_s^* = z^*$, $s \in \mathcal{S}$, implies that each limit point $((x_s^*, y_s^*)_{s \in \mathcal{S}}, z^*)$ is optimal for problem (9) and ω^* is optimal for the dual problem (8).

Thus, for all $t_{max} \geq 1$, we have shown that $\lim_{k \rightarrow \infty} (x_s^k - z^k) = 0$, $s \in \mathcal{S}$, and $\lim_{k \rightarrow \infty} \phi(\omega^k) = \zeta^{LD}$. By similar reasoning to that used in the $t_{max} \geq 2$ case, it is straightforward that, for all $\alpha \in \mathbb{R}$, we also have $\lim_{k \rightarrow \infty} \phi^k = \zeta^{LD}$. \square

While using a large value of t_{max} more closely matches Algorithm 3 to the original PH algorithm as described in Algorithm 1, we are motivated to use a small value of t_{max} since the work per iteration is proportional to t_{max} . Specifically, each iteration requires solving $t_{max}|\mathcal{S}|$ MILP subproblems and $t_{max}|\mathcal{S}|$ continuous convex quadratic subproblems. For reference, Algorithm 1 applied to problem (3) requires the solution of $|\mathcal{S}|$ MIQP subproblems for each ω update and $|\mathcal{S}|$ MILP subproblems for each Lagrangian bound ϕ computation.

3.2. Obtaining primal feasible solutions. As the FW-PH algorithm is designed to solve (9) (which is an optimization over $\text{conv}(K_s)$ rather than K_s), any solution it returns at convergence may not satisfy the integrality requirements and hence may not be primal feasible. Nevertheless, the information returned by FW-PH at termination may be exploited heuristically to derive primal feasible solutions. We suggest two simple heuristic strategies which use the solution $((x_s^k, y_s^k)_{s \in \mathcal{S}}, z^k, \omega^k)$

returned by FW-PH, as defined in Algorithm 3. These strategies may be used regardless of whether or not convergence has been achieved at termination. Both strategies take advantage of the assumption of relatively complete recourse: They *evaluate* a candidate first-stage solution by solving each of the $|\mathcal{S}|$ single-scenario problems with its first-stage variables fixed to the candidate values.

The first heuristic strategy, which we call H1, consists of evaluating each distinct solution in the set of solutions $\{\hat{x}_s^k : s \in \mathcal{S}\}$, obtained in the last execution of line 7 of Algorithm 3, as a candidate first-stage solution.

The second strategy, H2, consists of solving the MIQPs, one for each $s \in \mathcal{S}$, that would have been solved in PH (line 12 in Algorithm 1) using $z = z^k$, $\omega = \omega^k$, and considering $D_s = K_s$ for $s \in \mathcal{S}$, and evaluating each distinct first-stage solution found. Notice that either strategy may generate multiple candidate first-stage solutions, in particular when the FW-PH convergence criterion is not met at termination. In this case, the one evaluated to yield the best objective function value is selected. In section 4, we provide numerical results that assess the performance of these strategies.

4. Numerical experiments. We performed computations using a C++ implementation of Algorithm 1 ($D_s = K_s$, $s \in \mathcal{S}$) and Algorithm 3 using CPLEX 12.5 [18] as the solver for all subproblems. For reading SMPS files into scenario-specific subproblems and for their interface with CPLEX, we used modified versions of the COIN-OR [3] Smi and Osi libraries. The computing environment is the Raijin cluster maintained by Australia's National Computing Infrastructure (NCI) and supported by the Australian Government [1]. The Raijin cluster is a high-performance computing (HPC) environment which has 3592 nodes (system units), 57472 cores of Intel Xeon E5-2670 processors with up to 8 GB PC1600 memory per core (128 GB per node). All experiments were performed in a serial setting using a single node and one thread per CPLEX solve.

In the experiments with Algorithms 1 and 3, we set the convergence tolerance at $\epsilon = 10^{-3}$. For Algorithm 3, we set $t_{max} = 1$. Also, for all experiments performed, we set $\omega^0 = 0$. In this case, convergence of our algorithm requires that (17) holds, which can be guaranteed during the initialization of the inner approximations $(V_s^0)_{s \in \mathcal{S}}$. Under the assumption of relatively complete resource, a straightforward mechanism for ensuring that (17) holds is to solve the recourse problems for any fixed $\hat{x} \in X$. Specifically, for each $s \in \mathcal{S}$, let $\hat{y}_s \in \arg \min_y \{q_s^\top y : (\hat{x}, y) \in K_s\}$ and initialize V_s^0 for each $s \in \mathcal{S}$ so that $\{(\hat{x}, \hat{y}_s)\} \in V_s^0$. For the computational experiments, we take \hat{x} to be the first-stage variables from the solution to a single-scenario problem for one arbitrary scenario, say scenario $1 \in \mathcal{S}$, and enrich the sets V_s^0 by also including the solution to a single-scenario problem for s . In each case, the single-scenario problem is a Lagrangian problem of the form of (7), with $\omega_s := \omega_s^0$. Specifically, we initialize $V_1^0 := \{(x_1^0, y_1^0)\}$ and for each $s \in \mathcal{S}$, $s \neq 1$, initialize $V_s^0 := \{(x_s^0, y_s^0), (x_1^0, \bar{y}_s)\}$, where (x_s^0, y_s^0) solves $\min_{x,y} \{(c + \omega_s^0)^\top x + q_s^\top y : (x, y) \in K_s\}$ and \bar{y}_s solves $\min_y \{q_s^\top y : (x_1^0, y) \in K_s\}$ for each $s \in \mathcal{S}$.

Experiments were performed on eight instances of three distinct problems, namely the capacitated facility location problem (CAP) from [7], the dynamic capacity allocation problem (DCAP) available in [5], and the server location under uncertainty problem (SSLP), first introduced in [29].

The CAP problem is a two-stage SMIP with pure binary first-stage and continuous second-stage variables arising in the context of network design. We selected the instances coded as 101 and 102 in [7], using the first 250 from a list of 5000 scenarios available. The DCAP problem is a two-stage SMIP arising in dynamic capacity ac-

TABLE 1
Result summary for CAP problem instances: dual bounds.

ρ	Gap (%)			# Iterations			Time		
	PH	FW-PH		PH	FW-PH		PH	FW-PH	
		$\alpha = 0$	$\alpha = 1$		$\alpha = 0$	$\alpha = 1$		$\alpha = 0$	$\alpha = 1$
20	0.05	0.17	0.09	509	398	445	T	T	T
100	0.01	0.00	0.00	178	446	440	1975.91	T	T
500	0.07	0.00	0.00	540	92	93	T	931.84	986.83
1000	0.15	0.00	0.00	544	127	130	T	1345.04	1425.90
2500	0.34	0.00	0.00	581	259	274	T	3087.30	3276.03
5000	0.66	0.00	0.00	33	473	468	293.03	T	T
7500	0.99	0.00	0.00	28	18	19	225.66	138.80	170.14
15000	1.59	0.00	0.00	545	28	33	T	246.65	283.53

(a) CAP-101-250; absolute percentage gap based on the known optimal value 733827.32.

ρ	Gap (%)			# Iterations			Time		
	PH	FW-PH		PH	FW-PH		PH	FW-PH	
		$\alpha = 0$	$\alpha = 1$		$\alpha = 0$	$\alpha = 1$		$\alpha = 0$	$\alpha = 1$
20	0.47	0.46	0.49	422	426	412	T	T	T
100	0.01	0.00	0.00	219	408	405	3343.29	T	T
500	0.08	0.00	0.00	48	46	46	757.09	524.11	540.72
1000	0.13	0.00	0.00	24	25	24	297.34	271.72	286.68
2500	0.29	0.00	0.00	13	16	16	151.72	160.46	171.43
5000	0.61	0.00	0.00	14	18	18	156.90	170.86	188.87
7500	0.93	0.00	0.00	17	22	23	187.08	224.37	237.81
15000	1.91	0.00	0.00	22	39	42	228.26	450.64	436.41

(b) CAP-102-250; absolute percentage gap based on the known optimal value 788996.97.

quisition and allocation under uncertainty. The problem has mixed-integer first-stage variables and pure binary second-stage variables. We selected the instances coded as 233 and 243 (which encodes the number of resources, tasks, and periods, respectively), using all 500 scenarios available. The SSLP problem is a two-stage SMIP arising in server location under uncertainty. The problem has pure binary first-stage variables and mixed-binary second-stage variables. We considered the instances coded as 5-25-50, 5-25-100, 10-50-100, and 15-45-15 (which encode the number of servers, clients, and scenarios, respectively). Details concerning the mathematical formulation, available optimal values, and best known bounds for these problems are described in detail in [7] and [27] and are accessible at [2] for the DCAP and SSLP problems.

Two sets of Algorithm 3 experiments correspond to variants considering $\alpha = 0$ and $\alpha = 1$. For each problem, computations were performed for different penalty values $\rho > 0$. The penalty values used in the experiments for the SSLP instances were chosen to include those penalties that are tested in a computational experiment with PH whose results are depicted in Figure 2 of [15]. For the other problem instances, the set of penalty values ρ tested is chosen to capture a reasonably wide range of performance potential for both PH and FW-PH. All computational experiments were allowed to run for a maximum of two hours in wall clock time.

Tables 1–3 provide a summary indicating the quality of the Lagrangian bounds ϕ computed at the end of each experiment for the eight problems with varying penalty parameter ρ . In each of these tables, the first column lists the values of the penalty parameter ρ , while the following are presented for PH and FW-PH (for both $\alpha = 0$ and

TABLE 2
Result summary for DCAP problem instances: dual bounds.

ρ	Gap (%)			# Iterations			Time		
	PH	FW-PH		PH	FW-PH		PH	FW-PH	
		$\alpha = 0$	$\alpha = 1$		$\alpha = 0$	$\alpha = 1$		$\alpha = 0$	$\alpha = 1$
2	0.13	0.12	0.12	2234	576	570	T	T	T
5	0.22	0.09	0.09	2367	561	559	T	T	T
10	0.23	0.07	0.08	2583	592	573	T	T	T
20	0.35	0.07	0.07	2539	572	567	T	T	T
50	1.25	0.06	0.06	2721	578	580	T	T	T
100	1.29	0.06	0.06	2755	428	438	T	4016.29	4492.36
200	2.58	0.06	0.06	2667	256	262	T	1707.97	1848.49
500	2.58	0.07	0.07	2839	244	246	T	1799.88	1569.58

(a) *DCAP-233-500; absolute percentage gap based on the best known upper bound 1737.73.*

ρ	Gap (%)			# Iterations			Time		
	PH	FW-PH		PH	FW-PH		PH	FW-PH	
		$\alpha = 0$	$\alpha = 1$		$\alpha = 0$	$\alpha = 1$		$\alpha = 0$	$\alpha = 1$
2	0.14	0.18	0.18	1710	558	577	T	T	T
5	0.20	0.13	0.13	2108	570	562	T	T	T
10	0.29	0.11	0.11	2110	562	559	T	T	T
20	0.52	0.10	0.10	2233	570	577	T	T	T
50	0.70	0.10	0.10	2355	578	579	T	T	T
100	1.32	0.09	0.09	2504	393	395	T	3744.33	3849.53
200	1.40	0.10	0.09	2568	244	261	T	1866.03	1854.85
500	2.11	0.10	0.10	2486	180	165	T	983.41	884.66

(b) *DCAP-243-500; absolute percentage gap based on the known optimal value 2167.51.*

$\alpha = 1$) computations in the remaining columns: (1) the absolute percentage gap $|(\zeta^* - \phi)/\zeta^*| * 100\%$ between the computed Lagrangian bound ϕ and some reference value ζ^* that is either a known optimal value for the problem or a known best bound thereof (column “Gap (%)”); (2) the total number of dual updates (“# Iterations”); and (3) the indication of whether the algorithm terminated due to the time limit, indicated by letter “T”, or the satisfaction of the convergence criterion $\sqrt{\sum_{s \in \mathcal{S}} p_s \|x_s^k - z^{k-1}\|_2^2} < \epsilon$, indicated by the display of the time elapsed when convergence was attained (column “Time”).

The following observations can be made from the results presented in Tables 1–3. First, for small values of the penalty ρ , there is no clear preference between the bounds ϕ generated by PH and FW-PH. However, for higher penalties, the bounds ϕ obtained by FW-PH are consistently of better quality (i.e., higher) than those obtained by PH, regardless of the variant used (i.e., $\alpha = 0$ or $\alpha = 1$). This tendency is illustrated, for example, in Table 2(a), where the absolute percentage gap of the Lagrangian lower bound with the known optimal value was 0.06% with $\rho = 200$ for FW-PH ($\alpha = 0$), while it was 2.58% for the same value of ρ for PH. This improvement is consistently observed for the other problems and the other values of ρ that are not too close to zero. Also, FW-PH did not terminate with suboptimal convergence or display cycling behavior for any of the penalty values ρ in any of the problems considered. For example, all experiments considered in Table 3(a) terminated due to convergence.

The percentage gaps suggest that the convergence for PH was suboptimal, while it was optimal for FW-PH. Moreover, it is possible to see from these tables that the

TABLE 3
Result summary for SSLP problem instances: dual bounds.

ρ	Gap (%)			# Iterations			Time		
	PH	FW-PH		PH	FW-PH		PH	FW-PH	
		$\alpha = 0$	$\alpha = 1$		$\alpha = 0$	$\alpha = 1$		$\alpha = 0$	$\alpha = 1$
1	0.30	0.00	0.00	105	115	116	225.80	150.63	151.52
2	0.73	0.00	0.00	51	56	56	107.85	71.56	72.07
5	0.91	0.00	0.00	25	26	27	51.77	33.43	34.88
15	3.15	0.00	0.00	12	16	17	22.00	20.59	21.95
30	6.45	0.00	0.00	12	18	18	18.44	23.29	24.00
50	9.48	0.00	0.00	18	25	26	21.00	34.37	37.89
100	9.48	0.00	0.00	8	45	45	7.95	62.20	67.77

(a) *SSLP-5-25-50; absolute percentage gap based on the known optimal value* −121.60.

ρ	Gap (%)			# Iterations			Time		
	PH	FW-PH		PH	FW-PH		PH	FW-PH	
		$\alpha = 0$	$\alpha = 1$		$\alpha = 0$	$\alpha = 1$		$\alpha = 0$	$\alpha = 1$
1	0.16	0.00	0.00	82	97	90	385.08	266.05	248.92
2	0.45	0.00	0.00	42	43	44	196.76	119.57	121.30
5	1.06	0.00	0.00	18	21	22	83.66	58.29	61.62
15	2.96	0.00	0.00	13	15	16	51.40	42.50	46.35
30	6.21	0.00	0.00	19	24	23	56.58	70.47	64.26
50	7.91	0.00	0.00	3123	38	36	T	113.21	107.54
100	7.91	0.00	0.00	27	74	70	44.60	223.73	216.66

(b) *SSLP-5-25-100; absolute percentage gap based on the known optimal value* −127.37.

ρ	Gap (%)			# Iterations			Time		
	PH	FW-PH		PH	FW-PH		PH	FW-PH	
		$\alpha = 0$	$\alpha = 1$		$\alpha = 0$	$\alpha = 1$		$\alpha = 0$	$\alpha = 1$
1	0.57	0.22	0.22	130	234	234	T	T	T
2	0.63	0.03	0.03	131	226	227	T	T	T
5	1.00	0.00	0.00	104	218	219	4885.74	T	T
15	2.92	0.00	0.00	33	45	118	1012.11	1463.75	3949.99
30	4.63	0.00	0.00	18	21	22	413.28	618.52	619.85
50	4.63	0.00	0.00	11	26	27	202.47	759.83	756.59
100	4.63	0.00	0.00	9	43	45	106.76	1302.04	1271.27

(c) *SSLP-10-50-100; percentage gap based on the known optimal value* −354.19.

ρ	Gap (%)			# Iterations			Time		
	PH	FW-PH		PH	FW-PH		PH	FW-PH	
		$\alpha = 0$	$\alpha = 1$		$\alpha = 0$	$\alpha = 1$		$\alpha = 0$	$\alpha = 1$
1	2.85	2.15	2.17	224	304	300	T	T	T
2	2.21	1.00	1.00	193	272	272	T	T	T
5	1.21	0.01	0.03	181	180	178	7021.35	T	T
15	4.13	0.00	0.00	421	84	86	T	5022.34	4986.36
30	7.89	0.00	0.00	35	66	68	424.76	1873.24	1866.31
50	7.89	0.00	0.00	23	67	65	257.40	992.90	1020.19
100	7.89	0.00	0.00	6	69	62	32.25	562.65	428.18

(d) *SSLP-15-45-15; percentage gap based on the known optimal value* −253.60.

quality of the bounds ϕ obtained using FW-PH were not as sensitive to the value of the penalty parameter ρ as that obtained using PH.

The FW-PH with $\alpha = 0$ versus PH convergence profiles for the experiments performed are given in Figures 1–4, in which we provide plots of wall time versus Lagrangian bound values based on profiling of varying penalty for four of the eight problems considered. The times scales for the plots have been set such that trends are meaningfully depicted (1000s for CAP and DCAP instances, 100s for SSLP-5-25-50, and 7000s for SSLP-15-45-15). The trend of the Lagrangian bounds is depicted with solid lines for FW-PH with $\alpha = 0$ and with dashed lines for PH. Plots providing the same comparison for FW-PH with $\alpha = 1$ are provided in Appendix A.

As seen in the plots of Figures 1–4, the Lagrangian bounds ϕ generated with PH tend to converge suboptimally, often displaying cycling, for large penalty values. In terms of the quality of the bounds obtained, while there is no clear winner when low penalty ρ values are used, for large penalties, the quality of the bounds ϕ generated with FW-PH is consistently better than for the bounds generated with PH, regardless of the α value. This last observation is significant because the effective use of large penalty values ρ in methods based on augmented Lagrangian relaxation tends to yield the most rapid early iteration improvement in the Lagrangian bound; this point is most clearly illustrated in the plot of Figure 3. The remaining plots have been omitted due to space limitations.




FIG. 1. Convergence profile for CAP-101-250 (PH and FW-PH with $\alpha = 0$).

To assess the performance of the ideas discussed in section 3.2 concerning the generation of primal feasible solutions, we performed the following experiments. For PH, we used the solution of MIQPs (calculated in line 12 in Algorithm 1) returned by Algorithm 1 to give candidate first-stage solutions. Whenever PH has converged, a unique nonanticipative (first-stage) primal feasible integral solution is returned. Otherwise, PH might obtain distinct solutions for distinct scenarios; we evaluate all

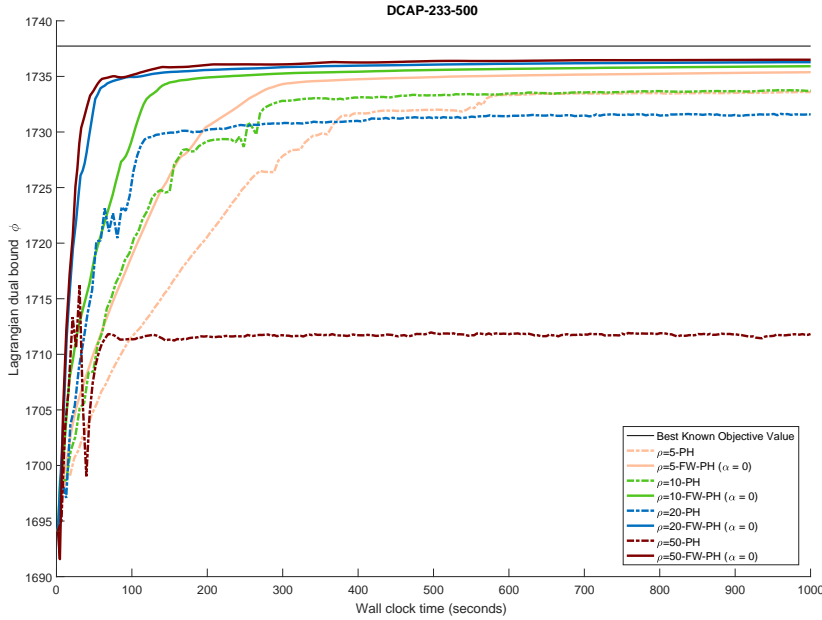


FIG. 2. Convergence profile for DCAP-233-500 (PH and FW-PH with $\alpha = 0$).

TABLE 4
Result summary for CAP problem instances: primal bounds.

ρ	Gap (%)			Time			ρ	Gap (%)			Time		
	PH	H1	H2	PH	H1	H2		PH	H1	H2	PH	H1	H2
20	0.00	0.00	0.00	130.26	984.17	288.24	20	0.00	0.05	0.00	551.42	1504.03	537.83
100	0.00	0.00	0.00	3.74	102.43	13.90	100	0.00	0.00	0.00	3.70	545.65	15.76
500	0.00	0.00	0.00	8.26	6.32	9.09	500	0.00	0.00	0.00	4.47	2.95	10.53
1000	0.00	0.00	0.00	7.89	2.97	9.19	1000	0.00	0.00	0.00	2.94	3.87	10.21
2500	0.00	0.00	0.00	19.42	5.44	12.02	2500	0.00	0.00	0.00	3.11	6.81	9.94
5000	0.00	0.00	0.00	3.36	7.98	13.27	5000	0.00	0.00	0.00	2.97	2.41	9.65
7500	0.00	0.00	0.00	3.29	2.64	8.28	7500	0.00	0.00	0.00	3.14	2.85	9.44
15000	0.10	0.00	0.00	8.83	3.10	8.06	15000	0.05	0.00	0.00	3.04	2.91	10.05

(a) CAP-101-250; absolute percentage gap based on the known optimal value 733827.32. (b) CAP-102-250; absolute percentage gap based on the known optimal value 788996.97.

distinct solutions and report that with the best objective function value. For FW-PH, we analyze the two distinct strategies discussed in section 3, referred to as H1 and H2, respectively.

In Tables 4–6, the first column lists the values of the penalty parameter ρ while the remaining columns present for PH and FW-PH (for H1 and H2): (1) the absolute percentage gap between ζ^* (i.e., either a known optimal value for the problem or a known best bound thereof) and the primal bound ζ obtained as described, which is given by $|(\zeta - \zeta^*)/\zeta^*|$, expressed as a percentage (column “Gap (%)”); and (2) the total wall clock time taken to evaluate all candidate solutions and select that with best objective function value. The results presented are for the variant of FW-PH considering $\alpha = 0$ only due to space limitations.




FIG. 3. Convergence profile for *SSLP-5-25-50* (PH and FW-PH with $\alpha = 0$).

TABLE 5
Result summary for DCAP problem instances: primal bounds.

		Gap (%)			Time		
ρ	PH	H1	H2	PH	H1	H2	
2	0.61	0.05	0.36	216.75	234.30	234.58	
5	0.52	1.15	0.50	207.56	311.18	204.83	
10	0.36	0.21	0.42	176.55	287.94	178.96	
20	0.01	0.06	0.41	131.18	229.91	162.12	
50	0.41	0.20	0.42	90.93	300.94	122.32	
100	0.42	0.41	0.42	65.22	300.71	92.03	
200	0.06	2.84	0.42	63.24	313.84	92.64	
500	0.42	0.12	0.42	40.19	298.39	65.04	

		Gap (%)			Time		
ρ	PH	H1	H2	PH	H1	H2	
2	0.05	0.17	0.01	290.91	405.02	278.85	
5	0.01	0.09	0.01	205.49	445.42	227.52	
10	0.05	0.16	0.01	173.84	407.50	186.12	
20	0.05	0.01	0.01	155.00	413.32	164.86	
50	0.01	0.46	0.01	149.15	439.40	145.58	
100	0.96	0.01	0.01	47.81	393.20	136.77	
200	1.05	0.01	0.01	49.28	448.63	136.00	
500	1.05	0.40	0.01	54.06	429.09	108.74	

(a) *DCAP-233-500*; absolute percentage gap based on the best known upper bound 1737.73. (b) *DCAP-243-500*; absolute percentage gap based on the known optimal value 2167.51.

As can be seen from these results, despite the simplicity of the proposed heuristics to obtain primal feasible solutions, in most cases the primal bounds generated are of good quality and often superior to those obtained from PH. Strategy H2 more often presents better performance when compared to H1, but with no clear winner between them in terms of percentage gap and time. The time taken for PH and FW-PH to evaluate the solutions is similar for the two methods, despite the heuristic used. One can notice that, in cases in which convergence is not observed (typically associated with the consideration of smaller penalty ρ), the time taken to evaluate solutions is typically higher due to the number of first-stage solutions to be considered. Overall, these results suggest that it is possible in practice to employ these heuristics as the times required are not prohibitive, particularly when convergence to a reasonably good dual bound has been observed.

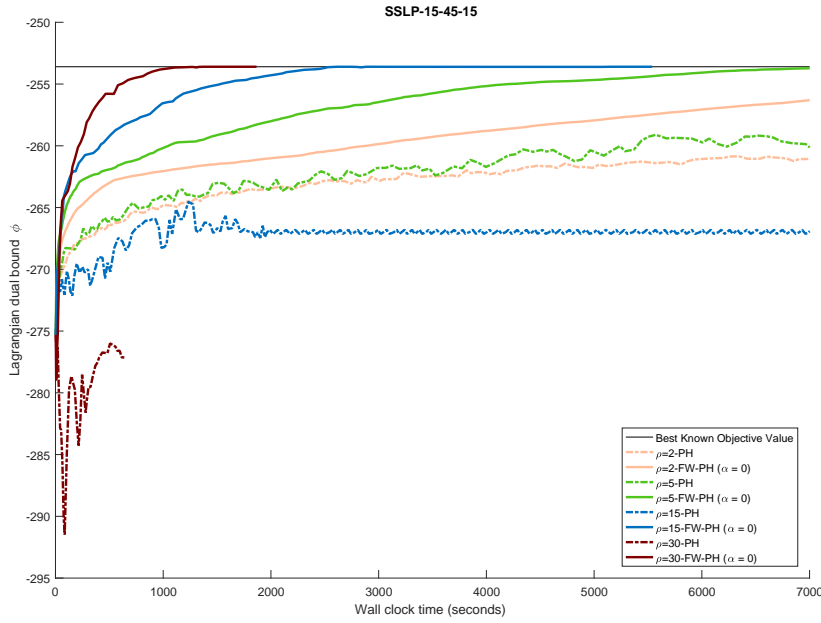


FIG. 4. Convergence profile for SSLP-15-45-15 (PH and FW-PH with $\alpha = 0$).

TABLE 6
Result summary for SSLP problem instances: primal bounds.

ρ	Gap (%)			Time		
	PH	H1	H2	PH	H1	H2
1	0.00	0.00	0.00	0.02	0.02	1.03
2	0.00	0.00	0.00	0.02	0.01	1.07
5	0.00	0.00	0.00	0.01	0.02	0.96
15	0.00	0.00	0.00	0.02	0.01	0.55
30	0.00	0.00	0.00	0.02	0.02	0.36
50	0.00	0.00	0.00	0.01	0.02	0.18
100	2.15	0.00	0.00	0.02	0.02	0.08

ρ	Gap (%)			Time		
	PH	H1	H2	PH	H1	H2
1	0.00	0.00	0.00	0.03	0.22	2.50
2	0.00	0.00	0.00	0.03	0.06	2.36
5	0.00	0.00	0.00	0.02	0.03	2.08
15	0.00	0.00	0.00	0.03	0.03	1.14
30	0.00	0.00	0.00	0.05	0.03	0.74
50	0.00	0.00	0.00	0.16	0.03	0.37
100	1.40	0.00	0.00	0.04	0.03	0.16

(a) SSLP-5-25-50; absolute percentage gap based on the known optimal value -121.60 .

(b) SSLP-5-25-100; absolute percentage gap based on the known optimal value -127.37 .

ρ	Gap (%)			Time		
	PH	H1	H2	PH	H1	H2
1	0.00	1.42	0.00	2.69	39.12	32.25
2	0.00	0.00	0.00	21.71	28.98	29.70
5	0.00	0.00	0.00	1.22	20.79	20.35
15	0.00	0.00	0.00	1.12	2.65	13.17
30	0.00	0.00	0.00	1.60	1.16	8.98
50	0.00	0.00	0.00	1.11	1.21	6.93
100	0.00	0.00	0.00	1.11	1.11	5.78

ρ	Gap (%)			Time		
	PH	H1	H2	PH	H1	H2
1	0.00	0.79	0.00	3.88	5.10	32.63
2	0.00	0.00	0.00	5.82	5.27	29.27
5	0.00	0.16	0.00	0.95	32.87	22.57
15	0.00	0.00	0.00	2.51	0.98	4.10
30	0.16	0.00	0.00	1.54	0.95	1.56
50	0.52	0.00	0.00	0.90	0.94	2.11
100	0.00	0.00	0.00	0.94	0.98	1.52

(c) SSLP-10-50-100; absolute percentage gap based on the known optimal value -354.19 .

(d) SSLP-15-45-15; absolute percentage gap based on the known optimal value -253.60 .

5. Conclusions and future work. In this paper, we have presented an alternative approach to computing nonanticipativity Lagrangian bounds associated with SMIPs that combines ideas from the progressive hedging (PH) and the Frank–Wolfe (FW) methods. We first note that while Lagrangian bounds can be recovered with PH, this requires—for each iteration and each scenario—the solution of an additional MILP subproblem in addition to the MIQP subproblem. Furthermore, when using the PH method directly, the Lagrangian bounds may converge suboptimally, cycle (for large penalties), or converge very slowly (for small penalties).

To overcome the lack of theoretical support for the above use of PH, we first described a straightforward integration of PH and an FW-like approach such as the simplicial decomposition method (SDM), where SDM is used to compute the primal updates in PH. Its convergence only requires noting that SDM applied to a convex problem with a bounded polyhedral constraint set terminates finitely with optimal convergence. However, for the stated goal of computing high-quality Lagrangian bounds efficiently, the benefits of relying on the optimal convergence of SDM is far outweighed by the computational costs incurred.

As an alternative, we propose the contributed algorithm, FW-PH, that is analyzed under general assumptions on how the Lagrangian bounds are computed and on the number of SDM iterations used for each dual update. Furthermore, under mild assumptions on the initialization of the algorithm, FW-PH only requires the solution of a MILP subproblem and a continuous convex quadratic subproblem for each iteration and each scenario. FW-PH is versatile enough to handle a wide range of SMIPs with integrality restrictions in any stage, while providing rapid improvement in the Lagrangian bound in the early iterations that is consistent across a wide range of penalty parameter values. Although we have opted to focus on two-stage problems with recourse, the generalization of the proposed approach to the multi-stage case is also possible.

The numerical results are encouraging as they suggest that the proposed FW-PH method applied to SMIP problems usually outperforms the traditional PH method with respect to how quickly the quality of the generated Lagrangian bounds improves. This is especially true with the use of larger penalty values. For all problems considered and for all but the smallest penalties considered, the FW-PH method displayed better performance over PH in terms of the quality of the final Lagrangian bounds at the end of the allotted wall clock time.

The improved performance of FW-PH over PH for large penalties is significant because it is the effective use of large penalties enabled by FW-PH that yields the most rapid initial dual improvement. This last feature of FW-PH would be most helpful in its use within a branch-and-bound or branch-and-cut framework for providing strong lower bounds (in the case of minimization). In addition to being another means to compute Lagrangian bounds, PH would still have a role in such frameworks as a heuristic for computing a primal feasible solution to the SMIP, thus providing (in the case of minimization) an upper bound on the optimal value. In fact, as demonstrated by our numerical experiments, straightforward combinations of both methods provide heuristics that are capable of generating very good primal feasible solutions for the problems considered. This suggests that the development of more sophisticated heuristics is also a promising avenue of research.

Possible directions for future research include the following. First, FW-PH inherits the potential for parallelization from PH. Experiments for exploring the benefit of parallelization are therefore warranted. Second, the theoretical support of FW-PH can be strengthened with a better understanding of the behavior of PH (and its gen-

eralization ADMM) applied to infeasible problems. Finally, FW-PH can benefit from a better understanding of how the proximal term penalty coefficient can be varied to improve performance.

Appendix A. Additional plots for PH vs. FW-PH for $\alpha = 1$.

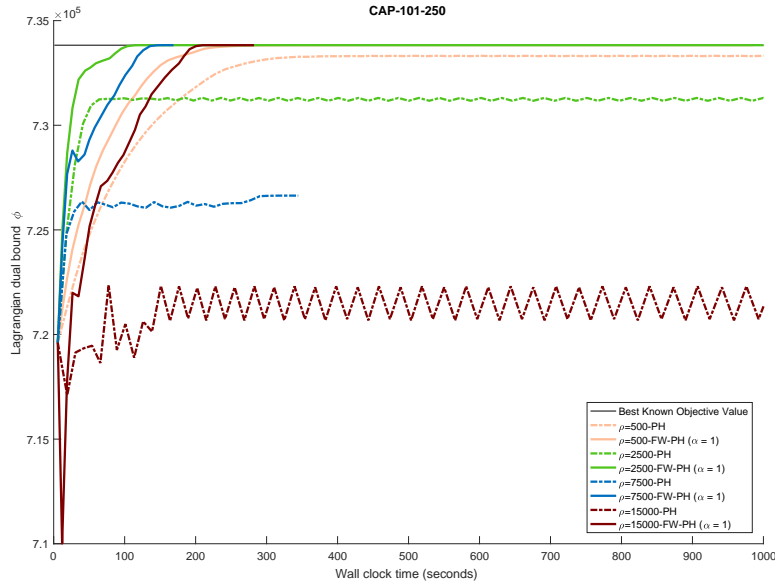


FIG. 5. Convergence profile for CAP-101-250 (PH and FW-PH with $\alpha = 1$).




FIG. 6. Convergence profile for DCAP-233-500 (PH and FW-PH with $\alpha = 1$).




FIG. 7. Convergence profile for *SSLP-5-25-50* (PH and FW-PH with $\alpha = 1$).




FIG. 8. Convergence profile for *SSLP-15-45-15* (PH and FW-PH with $\alpha = 1$).

REFERENCES

- [1] <http://www.nci.org.au>. Last accessed 20 April, 2016.
- [2] <http://www2.isye.gatech.edu/~sahmed/siplib/sslpl/sslpl.html>. Last accessed 23 December, 2015.
- [3] *COmputational Infrastructure for Operations Research*, <http://www.coin-or.org/>. Last accessed 28 January, 2016.
- [4] A. AGUSTÍN, A. ALONSO-AYUSO, L. F. ESCUDERO, AND C. PIZARRO, *On air traffic flow management with rerouting. Part II: Stochastic case*, European J. Oper. Res., 219 (2012), pp. 167–177.
- [5] S. AHMED AND R. GARCIA, *Dynamic capacity acquisition and assignment under uncertainty*, Ann. Oper. Res., 124 (2003), pp. 267–283.
- [6] F. BADILLA VELIZ, J.-P. WATSON, A. WEINTRAUB, R. J.-B. WETS, AND D. L. WOODRUFF, *Stochastic optimization models in forest planning: A progressive hedging solution approach*, Ann. Oper. Res., 232 (2015), pp. 259–274.
- [7] M. BODUR, S. DASH, O. GÜNLÜK, AND J. LUEDTKE, *Strengthened Benders Cuts for Stochastic Integer Programs with Continuous Recourse*, preprint, http://www.optimization-online.org/DB_FILE/2014/03/4263.pdf, 2014.
- [8] S. BOYD, N. PARIKH, E. CHU, B. PELEATO, AND J. ECKSTEIN, *Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers*, Foundations and Trends in Machine Learning, Vol. 3, Issue 1, Now Publishers, Hanover, MA, 2011, <http://dx.doi.org/10.1561/22000000016>.
- [9] C. C. CARØE AND R. SCHULTZ, *Dual decomposition in stochastic integer programming*, Oper. Res. Lett., 24 (1999), pp. 37–45.
- [10] J. CHEN, C. H. LIM, P. QIAN, J. LINDEROTH, AND S. J. WRIGHT, *Validating Sample Average Approximation Solutions with Negatively Dependent Batches*, preprint, <http://arxiv.org/abs/1404.7208>, 2014.
- [11] T. H. DE MELLO, *On rates of convergence for stochastic optimization problems under non-independent and identically distributed sampling*, SIAM J. Optim., 19 (2008), pp. 524–551.
- [12] J. ECKSTEIN AND D. BERTSEKAS, *On the Douglas–Rachford splitting method and the proximal point algorithm for maximal monotone operators*, Math. Program., 55 (1992), pp. 293–318.
- [13] M. FRANK AND P. WOLFE, *An algorithm for quadratic programming*, Naval Res. Logist. Quart., 3 (1956), pp. 149–154.
- [14] D. GABAY AND B. MERCIER, *A dual algorithm for the solution of nonlinear variational problems via finite element approximation*, Comput. Math. Appl., 2 (1976), pp. 17–40.
- [15] D. GADE, G. HACKEBEIL, S. M. RYAN, J.-P. WATSON, R. J.-B. WETS, AND D. L. WOODRUFF, *Obtaining lower bounds from the progressive hedging algorithm for stochastic mixed-integer programs*, Math. Program., 157 (2016), pp. 47–67.
- [16] M. R. HESTENES, *Multiplier and gradient methods*, J. Optim. Theory Appl., (1969), pp. 303–320.
- [17] C. HOLLOWAY, *An extension of the Frank and Wolfe method of feasible directions*, Math. Program., 6 (1974), pp. 14–27.
- [18] IBM CORPORATION, *IBM ILOG CPLEX V12.5*; available at <http://www-01.ibm.com/software/commerce/optimization/cplex-optimizer/>. Last accessed 28 Jan 2016.
- [19] A. KLEYWEGT, A. SHAPIRO, AND T. HOMEM-DE-MELLO, *The sample average approximation method for stochastic discrete optimization*, SIAM J. Optim., 12 (2001), pp. 479–502.
- [20] G. LAPORTE, F. V. LOUVEAUX, AND H. MERCURE, *The vehicle routing problem with stochastic travel times*, Transp. Sci., 26 (1992), pp. 161–170.
- [21] A. LØKKETANGEN AND D. WOODRUFF, *Progressive hedging and tabu search applied to mixed integer (0,1) multi-stage stochastic programming*, J. Heuristics, 2 (1996), pp. 111–128.
- [22] F. V. LOUVEAUX, *Discrete stochastic location models*, Ann. Oper. Res., 6 (1986), pp. 23–34.
- [23] M. LUBIN, K. MARTIN, C. PETRA, AND B. SANDIKÇI, *On parallelizing dual decomposition in stochastic integer programming*, Oper. Res. Lett., 41 (2013), pp. 252–258.
- [24] W. K. MAK, D. P. MORTON, AND R. K. WOOD, *Monte Carlo bounding techniques for determining solution quality in stochastic programs*, Oper. Res. Lett., 24 (1999), pp. 47–56.
- [25] G. NEMHAUSER AND L. WOLSEY, *Integer and combinatorial optimization*, Wiley-Intersci. Ser. Discrete Math. Optim., Wiley-Interscience, New York, 1988.
- [26] M. P. NOWAK AND W. RÖMISCH, *Stochastic Lagrangian relaxation applied to power scheduling in a hydro-thermal system under uncertainty*, Ann. Oper. Res., 100 (2000), pp. 251–272.
- [27] L. NTAIMO, *Decomposition Algorithms for Stochastic Combinatorial Optimization: Computational Experiments and Extensions*, Ph.D. thesis, University of Arizona, AZ, 2004.
- [28] L. NTAIMO, J. A. G. ARRUBLA, C. STRIPLING, J. YOUNG, AND T. SPENCER, *A stochastic*

- programming standard response model for wildfire initial attack planning*, Can. J. For. Res., 42 (2012), pp. 987–1001.
- [29] L. NTAIMO AND S. SEN, *The million-variable “march” for stochastic combinatorial optimization*, J. Global Optim., 32, pp. 385–400.
- [30] O. Y. ÖZALTIN, O. A. PROKOPYEV, A. J. SCHAEFER, AND M. S. ROBERTS, *Optimizing the societal benefits of the annual influenza vaccine: A stochastic programming approach*, Oper. Res., 59 (2011), pp. 1131–1143.
- [31] M. J. D. POWELL, *A method for nonlinear constraints in minimization problems*, in Optimization, R. Fletcher, ed., Academic Press, New York, 1969.
- [32] R. T. ROCKAFELLAR AND R. J.-B. WETS, *Scenarios and policy aggregation in optimization under uncertainty*, Math. Oper. Res., 16 (1991), pp. 119–147.
- [33] A. RUSZCZYŃSKI, *Nonlinear Optimization*, Princeton University Press, Princeton, NJ, 2006.
- [34] J. SALMERON, R. K. WOOD, AND D. P. MORTON, *A stochastic program for optimizing military sealift subject to attack*, Military Oper. Res., 14 (2009), pp. 19–39.
- [35] N. SHOR, *Minimization Methods for Non-Differentiable Functions*, Springer-Verlag, New York, 1985.
- [36] S. TAKRITI, J. BIRGE, AND E. LONG, *A stochastic model for the unit commitment problem*, IEEE Trans. Power Syst., 11 (1996), pp. 1497–1508.
- [37] M. W. TANNER, L. SATTENSPIEL, AND L. NTAIMO, *Finding optimal vaccination strategies under parameter uncertainty using stochastic programming*, Math. Biosci., 215 (2008), pp. 144–151.
- [38] B. VON HOHENBALKEN, *Simplicial decomposition in nonlinear programming algorithms*, Math. Program., 13 (1977), pp. 49–68.
- [39] J. WATSON AND D. WOODRUFF, *Progressive hedging innovations for a class of stochastic resource allocation problems*, Comput. Manag. Sci., 8 (2011), pp. 355–370.