

1 **An Asynchronous Bundle-Trust-Region Method for Dual Decomposition**
2 **of Stochastic Mixed-Integer Programming** *

3 KIBA EK KIM[†], COSMIN G. PETRA[‡], AND VICTOR M. ZAVALA[§]

4 **Abstract.** We present an asynchronous bundle-trust-region algorithm within the context of
5 Lagrangian dual decomposition for stochastic mixed-integer programs. The approach solves the La-
6 grangian master problem by using a bundle method with a trust-region constraint. This scheme en-
7 ables asynchronous computations and can thus help mitigate severe load imbalance issues (associated
8 with the solution of scenario subproblems) and improve parallel efficiency. We provide a convergence
9 analysis and an implementation of the proposed scheme. We also present extensive numerical results
10 on eighty instances of a large-scale stochastic unit commitment problem, and demonstrate that the
11 proposed approach provides significant reductions in solution time and achieves strong scaling.

12 **Key words.** stochastic integer programming, bundle methods, asynchronous, parallel comput-
13 ing

14 **AMS subject classifications.** 49M27, 65K05, 68W10, 90C10, 90C15

15 **1. Introduction.** We consider the Lagrangian dual decomposition (DD) of two-
16 stage stochastic mixed-integer programs (SMIPs) of the form

17 (1a)
$$\min_{x_j, y_j} \sum_{j=1}^N p_j (c^T x_j + q_j^T y_j)$$

18 (1b)
$$\text{s.t.} \quad \sum_{j=1}^N H_j x_j = 0,$$

19 (1c)
$$(x_j, y_j) \in G_j, \quad j \in J.$$

21 Here, $x_j \in \mathbb{R}^{n_1}$ and $y_j \in \mathbb{R}^{n_2}$ are decision variables associated to scenario $j \in J :=$
22 $\{1, \dots, N\}$. The matrices H_j are such that (1b) represent the *nonanticipativity* con-
23 straints $x_1 = x_2 = \dots = x_N$. We define the feasible sets:

24
$$G_j := \{(x, y) : Ax \geq b, T_j x + W_j y \geq h_j, x \in X, y \in Y\}, \quad j \in J,$$

26 and assume that these sets are bounded and nonempty. We also define mixed-integer
27 sets for x_j and y_j of the form $X := \mathbb{R}^{n_1-p_1} \times \mathbb{Z}^{p_1}$ and $Y := \mathbb{R}^{n_2-p_2} \times \mathbb{Z}^{p_2}$, respectively.

28 The DD of (1) is obtained by applying a Lagrangian relaxation of the nonantici-
29 pativity constraints (1b). The Lagrangian dual function is given by

30 (2)
$$D(\lambda) := \min_{x_j, y_j} \left\{ \sum_{j=1}^N L_j(x_j, y_j, \lambda) : (x_j, y_j) \in G_j, j \in J \right\},$$

31

32 where $L_j(x_j, y_j, \lambda) := p_j (c^T x_j + q_j^T y_j) + \lambda^T H_j x_j$ and λ are the dual variables of
33 the nonanticipativity constraints. The evaluation $D(\lambda)$ involves the solution of $|J|$

*Preprint number: ANL/MCS-8046-0917; IM release number: LLNL-JRNL-738242

[†]Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Avenue, Lemont, IL 60439, USA (kimk@anl.gov).

[‡]Lawrence Livermore National Laboratory, 7000 East Avenue, Livermore, CA 94550, USA (petra1@llnl.gov).

[§]Department of Chemical and Biological Engineering, University of Wisconsin-Madison, 1415 Engineering Dr, Madison, WI 53706, USA (victor.zavala@wisc.edu).

34 decoupled scenario subproblems. Consequently, the Lagrangian dual problem can be
 35 written as

$$36 \quad (3) \quad z_{LD} := \max_{\lambda} \sum_{j=1}^N D_j(\lambda),$$

37
 38 where,

$$39 \quad (4) \quad D_j(\lambda) := \min_{x_j, y_j} \{L_j(x_j, y_j, \lambda) : (x_j, y_j) \in G_j\}.$$

40
 41 We use λ^* to denote an optimal dual variable (i.e., satisfying $D(\lambda^*) = z_{LD}$).

42 We recall that the function $D_j(\lambda)$ is a piecewise concave function of λ . Moreover,
 43 $D_j(\cdot)$ is composed of a finite number of segments, which is no more than the number
 44 of extreme points of $\text{conv}(G_j)$ (the convex hull of G_j).

45 DD methods for SMIP problems have been widely studied in the literature (e.g.,
 46 [5, 1, 14]). These methods offer the ability to address integer recourse variables (i.e.,
 47 integer variables in the second stage and beyond) and the opportunity to evaluate the
 48 dual function by solving parallel subproblems. Because each scenario subproblem of
 49 the decomposition method is a mixed-integer program, however, parallel computations
 50 may suffer from significant load imbalance and decreased parallel efficiency.

51 A few studies reported in the literature have applied asynchronous schemes for
 52 DD. Ryan et al. [23] extended a (synchronous) scenario decomposition method pro-
 53 posed in [1] to conduct asynchronous computation of scenario subproblems. This
 54 method eliminates discrete feasible solutions sequentially instead of finding the best
 55 dual variable. Aravena and Papavasiliou [2] use an incremental subgradient search
 56 that updates dual variables based only on a subset of the scenario subproblem func-
 57 tions. A limitation of this approach is that subgradient methods rely on an appropriate
 58 choice of the step size, which is not straightforward to determine in practice.

59 Incremental methods that enable asynchronous computations have been widely
 60 studied in the context of large-scale convex optimization. The proposed methods
 61 are based mainly on subgradient-based search strategies [19, 15]. An asynchronous
 62 subgradient variant was proposed and analyzed in [20]. Bertsekas further extended
 63 the incremental subgradient method to a proximal variant [3] and an augmented
 64 Lagrangian variant [4]. Gaudioso et al. [10] studies the case of minimizing a function
 65 defined as the pointwise maximum over a number of convex functions, where only
 66 a subset of component functions are evaluated. As pointed out in [7], the method
 67 in [10] cannot be applied sums of functions (as needed in the context of DD). Emiel
 68 and Sagastizábal [7] consider the case of minimizing the sum of convex functions with
 69 focus on inexact function evaluations [16]. In recent work [24], an incremental bundle
 70 method was proposed to use lower and upper models for updating the stability center.
 71 This work also considers inexact function evaluations for the bundle model. Numerical
 72 performance for the schemes proposed in [10, 7, 24] was limited. Furthermore, none of
 73 these works address issues associated with parallel implementation and performance
 74 issues (such as load imbalancing).

75 An asynchronous implementation of a proximal bundle method is provided in [8],
 76 where bundle solutions are *implicitly* regularized in the objective function. Fischer and
 77 Helmberg [8] present a proximal bundle framework that can asynchronously choose
 78 and solve the subspace of a general convex function. Unfortunately, this framework
 79 suffers of important scalability limitations. In particular, the method stores and
 80 frequently accesses global data, which must be distributed for large-scale problems.

81 At every iteration, a significant amount of communication overhead would occur. As
 82 a result, computational performance and efficiency of this method remain unanswered
 83 [8].

84 Trust-region constraints provides an *explicit* regularization mechanism to develop
 85 bundle methods, as compared with the *implicit* regularization provided by proximal
 86 bundle methods. Linderoth and Wright [17] applied a parallel bundle-trust-region
 87 method and its asynchronous variant to the L-shaped method (also known as Ben-
 88 ders decomposition) for small- and medium-sized stochastic linear programming prob-
 89 lems [17]. A generalization of bundle methods with a unified form of penalty-like and
 90 trust-region-like stabilizing terms has been considered and analyzed in [9]. In this
 91 work, we further develop these ideas further to develop an incremental bundle meth-
 92 ods in the context of Lagrangian dual decomposition for solving large-scale SMIP
 93 problems. We prove that our incremental bundle method is convergent for any work-
 94 allocation policy (static against dynamic in Subsection 3.3), any trial-point selection
 95 policy (first-in-first-out against last-in-first-out in Subsection 3.1), any trust region
 96 norm (including 2-norm and ∞ -norms), and any bundle management step. We also
 97 perform extensive numerical experiments to demonstrate that the asynchronous im-
 98 plementation achieves significant improvements in parallel efficiency and solution time
 99 over a standard synchronous implementation.

100 The remainder of the paper is organized as follows. In Section 2 we present the
 101 bundle-trust-region method for Lagrangian DD and associated convergence analysis.
 102 Section 3 presents an asynchronous variant of the method and associated conver-
 103 gence analysis. In Section 4 we present the implementation of the methods in the
 104 open-source software package DSP [14] and numerical experiments. In Section 5, we
 105 summarize the paper and discuss possible directions for future research.

106 **2. A Bundle-Trust-Region Algorithm.** We propose a bundle-trust-region
 107 (BTR) algorithm to solve the Lagrangian dual problem (3). The proposed BTR
 108 method iteratively approximates the Lagrangian dual function $D(\lambda)$ by adding a set
 109 of cutting planes (cuts). We let $k \in \mathbb{Z}_+$ and $l \in \mathbb{Z}_+$ be the indices for major and minor
 110 iterations, respectively. Every major iteration updates the best lower bound, whereas
 111 every minor iterate updates the approximation of the Lagrangian dual function. We
 112 define the model function:

$$113 \quad (5a) \quad m_{k,l}(\lambda) := \max_{j \in J} \theta_j$$

$$114 \quad (5b) \quad \text{s.t. } \theta_j \leq D_j(\lambda^i) + (H_j x_j^i)^T (\lambda - \lambda^i), \quad i \in \mathcal{B}^{k,l}, \quad j \in J.$$

116 We recall that $H_j x_j^i \in \partial D_j(\lambda^i)$ for $j \in J$ and $\mathcal{B}^{k,l}$ is a set of cut indices at iteration
 117 (k, l) . We also recall that the model function $m_{k,l}(\cdot)$ outer-approximates $D(\cdot)$ (i.e.,
 118 $m_{k,l}(\lambda) \geq D(\lambda)$ for all $\lambda \in \mathbb{R}^{N_{n_1}}$ and (k, l)).

119 At iteration (k, l) , the master problem is given by

$$120 \quad (6a) \quad \max_{\lambda \in \mathbb{R}^{N_{n_1}}} m_{k,l}(\lambda)$$

$$121 \quad (6b) \quad \text{s.t. } \|\lambda - \lambda^k\| \leq \Delta_{k,l},$$

123 where λ^k is the TR center and $\Delta_{k,l} > 0$ is the TR size. The master problem (6) finds
 124 a new trial point $\lambda^{k,l}$ to evaluate $D(\cdot)$. We define the *predicted increase* of $D(\cdot)$ as:

$$125 \quad (7) \quad v_{k,l} := m_{k,l}(\lambda^{k,l}) - D(\lambda^k).$$

127 The TR center is updated as $\lambda^{k+1} \leftarrow \lambda^{k,l}$ if the sufficient decrease condition

$$128 \quad (8) \quad D(\lambda^{k,l}) \geq D(\lambda^k) + \xi v_{k,l},$$

130 is satisfied, where $\xi \in (0, 1/2)$. We call this type of iteration a *serious step*. In this
131 case, the master problem (6) is resolved with the updated TR center. Otherwise,
132 a new set of cuts(5b) is added to improve the model $m_{k,l}(\lambda)$. We call this type of
133 iteration a *null step*. The method terminates whenever

$$134 \quad (9) \quad v_{k,l} \leq \epsilon \cdot (1 + |D(\lambda^k)|)$$

136 holds for some $\epsilon \in \mathbb{R}_+$.

137 The TR size $\Delta_{k,l}$ need to be carefully updated to accelerate performance. For
138 example, if the TR size is too large, a number of null steps must be taken before each
139 serious step is taken. On the other hand, if the TR size is too small, the algorithm
140 takes serious steps at almost all iterations with only marginal improvements in the
141 lower bound. We thus devise tests for detecting whether the TR size is too large or
142 too small. To do so, we define the model approximation error at iterate (k, l) as:

$$143 \quad (10) \quad \delta^{k,l} := \sum_{j \in J} \delta_j^{k,l},$$

145 where,

$$146 \quad (11) \quad \delta_j^{k,l} := D_j(\lambda^{k,l}) + (H_j x_j^{k,l})^T (\lambda^k - \lambda^{k,l}) - D_j(\lambda^k).$$

148 By construction, $\delta_j^{k,l} \geq 0$ holds. We define the maximum model variation as:

$$149 \quad (12) \quad V_k := \max_{\lambda} \{D(\lambda) : \|\lambda - \lambda^k\| \leq 1\} - D(\lambda^k).$$

151 We deem the TR size too large if either $D(\lambda^k) - D(\lambda^{k,l})$ or the approximation error
152 $\delta^{k,l}$ are much larger than V_k . By construction, we have that:

$$153 \quad (13) \quad V_k \leq \max_{\lambda} \{m_{k,l}(\lambda) : \|\lambda - \lambda^k\| \leq 1\} - D(\lambda^k)$$

$$154 \quad (14) \quad \leq \frac{v_{k,l}}{\min\{1, \|\lambda^{k,l} - \lambda^k\|\}}.$$

156 Consequently, it is sufficient to test whether

$$157 \quad (15) \quad \max\{D(\lambda^k) - D(\lambda^{k,l}), \delta^{k,l}\} > \frac{v_{k,l}}{\min\{1, \|\lambda^{k,l} - \lambda^k\|\}}$$

159 holds in order to determine whether the TR size is too large. We now let

$$160 \quad (16) \quad \rho := \min\{1, \|\lambda^{k,l} - \lambda^k\|\} \frac{\max\{D(\lambda^k) - D(\lambda^{k,l}), \delta^{k,l}\}}{v_{k,l}}$$

162 and let τ^k count the iterations in which the TR size is not reduced. We then update
163 the TR size as follows:

- 164 1. If $\rho > 0$, then $\tau^{k+1} := \tau^k + 1$;
- 165 2. If $\rho > \bar{\rho}$ or $(\rho \in (0, \bar{\rho})$ and $\tau \geq \bar{\rho})$, then set $\tau^{k+1} := 0$, and

$$166 \quad (17) \quad \Delta_{k,l+1} \leftarrow \max\left\{\frac{\Delta_{k,l}}{\min\{\rho, \underline{\rho}\}}, \underline{\Delta}\right\}.$$

167

168 We deem the TR to be too small if a larger (i.e., better) Lagrangian function value is
 169 found and if the solution is bounded by the TR constraint. That is, whenever

$$170 \quad (18) \quad D(\lambda^{k+1}) \geq D(\lambda^k) + \frac{1}{2} v_{k,l} \quad \text{and} \quad \|\lambda - \lambda^k\| \leq \Delta_{k,l},$$

172 holds, we increase the TR size as

$$173 \quad (19) \quad \Delta_{k,l+1} \leftarrow \min\{2\Delta_{k,l}, \bar{\Delta}\}.$$

Algorithm 1 Bundle-Trust-Region Method

- 1: Initialize $\lambda^0 \in \mathbb{R}^{Nn_1}$, $\Delta_{0,0} \in [\underline{\Delta}, \bar{\Delta}]$, $\xi \in (0, 1/2)$, $\epsilon \geq 0$, $\mathcal{B}^{0,0} \leftarrow \{0\}$, $k \leftarrow 0$, and $l \leftarrow 0$.
 - 2: Solve the Lagrangian subproblem (4) to find $D_j(\lambda^0)$ and x_j^0 for all $j \in J$.
 - 3: Initialize the model function $m_{0,0}$.
 - 4: **loop**
 - 5: Solve the master (6) to $\lambda^{k,l}$.
 - 6: Solve the Lagrangian subproblem (4) to find $D_j(\lambda^{k,l})$ and $x_j^{k,l}$ for all $j \in J$.
 - 7: **if** $m_{k,l}(\lambda^{k,l}) - D(\lambda^k) \leq \epsilon(1 + |D(\lambda^k)|)$ **then** ▷ Termination test
 - 8: Stop
 - 9: **end if**
 - 10: **if** $D(\lambda^{k,l}) \geq D(\lambda^k) + \xi[m_{k,l}(\lambda^{k,l}) - D(\lambda^k)]$ **then** ▷ Serious step
 - 11: $\lambda^{k+1} \leftarrow \lambda^{k,l}$.
 - 12: Choose $\Delta_{k+1,0} \in [\Delta_{k,l}, \bar{\Delta}]$.
 - 13: $m_{k+1,0} \leftarrow m_{k,l}$.
 - 14: $k \leftarrow k + 1$ and $l \leftarrow 0$.
 - 15: **else** ▷ Null step
 - 16: Choose $\Delta_{k,l+1} \in [\underline{\Delta}, \Delta_{k,l}]$.
 - 17: Update $m_{k,l+1}$ by adding cuts (5b).
 - 18: $l \leftarrow l + 1$.
 - 19: **end if**
 - 20: **end loop**
-

175 **2.1. Algorithmic Steps.** We now summarize the steps of Algorithm 1. The
 176 BTR algorithm is initialized using the dual λ^0 and parameters $\Delta_{k,l}$, ξ , and ϵ (line 1).
 177 For a given λ^0 , the Lagrangian dual function is evaluated by solving the subprob-
 178 lem (4) for each scenario $j \in J$ (line 2). The model function $m_{0,0}$ is initialized by
 179 adding cuts (5b) generated at x_j^0 (line 3). The algorithm continues by finding a new
 180 dual value $\lambda^{k,l}$ (line 5), solving the Lagrangian subproblems (line 6), updating the
 181 TR (lines 11, 12, and 16), and updating the master problem (lines 13 and 17) until
 182 the termination criterion is met (line 8).

183 The BTR algorithm solves the Lagrangian dual problem and thus only provides
 184 a lower bound for the original SMIP (1). An upper bound for SMIP (1) can be
 185 obtained by evaluating first-stage variable solutions x_j^k obtained for each Lagrangian
 186 subproblem $D_j(\lambda^k)$ at iteration k . We note that at most $|J|$ first-stage solutions can
 187 be obtained in each iteration. The evaluation of first-stage solution x^k solves the
 188 second-stage recourse function $\mathcal{Q}(x^k) := \sum_{j=1}^N p_j Q_j(x^k)$, where

$$189 \quad (20) \quad Q_j(x) := \min_{y \in Y} \{q_j^T y : W_j y \geq h_j - T_j x\}.$$

190

191 We assume that $Q_j(x) = \infty$ if there does not exist recourse $y \in Y$ such that $W_j y \geq$
 192 $h_j - T_j x$ (i.e., the candidate x is infeasible). We recall that obtaining an optimal
 193 upper bound requires an exhaustive branch-and-bound scheme [5, 12].

194 **2.2. Convergence Analysis.** We now present a convergence analysis for **Algo-**
 195 **rithm 1**. We assume that $\epsilon = 0$ throughout this section. We first show that only a
 196 finite number of cuts (5b) are available to construct the model function $m_{k,l}(\cdot)$.

197 **LEMMA 2.1.** *Algorithm 1 can generate only a finite number of cuts (5b).*

198 *Proof.* The subgradients $(H_j x_j^k)$ at major iteration k are obtained at the vertices
 199 (x_j, y_j) of the polyhedra $\text{conv}(G_j)$, $j \in J$. Such vertices are finitely many in $\text{conv}(G_j)$
 200 for all $j \in J$. \square

201 We now show that the algorithm does not perform an infinite number of null steps.

202 **LEMMA 2.2.** *Suppose that Algorithm 1 takes a null step at iteration $(k, 0)$. There*
 203 *exists a finite $l > 0$ such that the algorithm either takes a serious step or terminates*
 204 *at iteration l .*

205 *Proof.* To establish a contradiction, suppose that the algorithm takes infinitely
 206 many null steps at major iteration k . Then, we have for $l \geq 0$

$$207 \quad (21) \quad m_{k,l}(\lambda^{k,l}) - D(\lambda^k) > \epsilon(1 + |D(\lambda^k)|) = 0$$

209 and

$$210 \quad (22) \quad D(\lambda^{k,l}) < D(\lambda^k) + \xi [m_{k,l}(\lambda^{k,l}) - D(\lambda^k)].$$

212 Let $\theta_j^{k,l}$ be the solution of the model $m_{k,l}(\lambda^{k,l})$. There exists some $j \in J$ such that
 213 linear inequalities generated at iteration (k, l) ,

$$214 \quad (23) \quad \theta_j \leq D_j(\lambda^{k,l}) + (H_j x_j^{k,l})^T (\lambda - \lambda^{k,l}),$$

216 are violated at $(\theta^{k,l}, \lambda^{k,l})$; because otherwise

$$217 \quad \sum_{j \in J} \theta_j^{k,l} - D(\lambda^{k,l}) = m_{k,l}(\lambda^{k,l}) - D(\lambda^{k,l})$$

$$218 \quad > m_{k,l}(\lambda^{k,l}) - D(\lambda^k) - \xi [m_{k,l}(\lambda^{k,l}) - D(\lambda^k)]$$

$$219 \quad > 0.$$

221 Here, the first and second inequalities hold due to (22) and (21), respectively. By
 222 **Lemma 2.1**, we have that either (21) or (22) will be violated after a finite number of
 223 null steps, contradicting the assumption. \square

224 We adapt the approach that obtains the lower bound of the predicted increase,
 225 as shown in [17].

226 **LEMMA 2.3.** *For $k \geq 0$ and $l \geq 0$, we have that*

$$227 \quad (24) \quad m_{k,l}(\lambda^{k,l}) - m_{k,l}(\lambda^k) \geq \min \left\{ \frac{\Delta_{k,l}}{\|\lambda^* - \lambda^k\|}, 1 \right\} [D(\lambda^*) - D(\lambda^k)].$$

228

229 *Proof.* We need only to show that (24) holds for the ∞ -norm, because $\|\cdot\|_\infty \leq$
 230 $\|\cdot\|_p \leq \|\cdot\|_r$ for $p > r > 0$. Suppose that we obtain an optimal step size $\alpha_{k,l}$ of the
 231 form

$$232 \quad \alpha_{k,l} := \arg \max_{\alpha \in [0,1]} \{m_{k,l}(\lambda^k + \alpha[\lambda^* - \lambda^k]) : \|\alpha[\lambda^* - \lambda^k]\| \leq \Delta_{k,l}\}.$$

234 We have

$$235 \quad \begin{aligned} m_{k,l}(\lambda^{k,l}) &\geq m_{k,l}(\lambda^k + \alpha_{k,l}[\lambda^* - \lambda^k]) \\ 236 \quad &\geq D(\lambda^k + \alpha_{k,l}[\lambda^* - \lambda^k]) \\ 237 \quad &\geq D(\lambda^k) + \alpha_{k,l}[D(\lambda^*) - D(\lambda^k)], \end{aligned}$$

239 where the first inequality holds because $\lambda^{k,l}$ is the maximizer of $m_{k,l}$, the second
 240 inequality holds because $m_{k,l}$ is an outer approximation of $D(\cdot)$, and the last inequality
 241 holds because of concavity of $D(\cdot)$. Since $m_{k,l}(\lambda^k) = D(\lambda^k)$, we have that

$$242 \quad (25) \quad m_{k,l}(\lambda^{k,l}) - m_{k,l}(\lambda^k) \geq m_{k,l}(\lambda^{k,l}) - D(\lambda^k) \geq \alpha_{k,l}[D(\lambda^*) - D(\lambda^k)].$$

244 Moreover, since $\|\alpha[\lambda^* - \lambda^k]\| \leq \Delta_{k,l}$, the optimal step size is given by

$$245 \quad (26) \quad \alpha_{k,l} = \min \left\{ \frac{\Delta_{k,l}}{\|\lambda^* - \lambda^k\|}, 1 \right\}.$$

247 Therefore, (24) is obtained from (25) and (26). \square

248 We now show that Algorithm 1 finds an optimal Lagrangian dual bound.

249 **THEOREM 2.4.** *Algorithm 1 delivers a dual iterate sequence $\{\lambda_k\}$ satisfying*
 250 $\lim_{k \rightarrow \infty} D(\lambda^k) \rightarrow D(\lambda^*)$.

251 *Proof.* Let l_k be the iteration index in which λ^{k,l_k} takes a serious step and thus
 252 $\lambda^{k+1} = \lambda^{k,l_k}$. Since

$$253 \quad (27) \quad m_{k,l_k}(\lambda^{k,l_k}) - D(\lambda^k) \geq \epsilon(1 + |D(\lambda^k)|) > 0$$

255 and

$$256 \quad (28) \quad D(\lambda^{k,l_k}) - D(\lambda^k) \geq \xi[m_{k,l_k}(\lambda^{k,l_k}) - D(\lambda^k)] > 0,$$

258 we have that $D(\lambda^{k+1}) - D(\lambda^k) > 0$. Since $\{D(\lambda^k)\}$ is an increasing sequence bounded
 259 above by z_{LD} , we have that

$$260 \quad (29) \quad \lim_{k \rightarrow \infty} D(\lambda^{k+1}) - D(\lambda^k) = 0.$$

262 Moreover, from Lemma 2.3 we have that

$$263 \quad \begin{aligned} m_{k,l_k}(\lambda^{k,l_k}) - m_{k,l_k}(\lambda^k) &= D(\lambda^{k+1}) - D(\lambda^k) \\ 264 \quad &\geq \min \left\{ \frac{\Delta_{k,l_k}}{\|\lambda^* - \lambda^k\|}, 1 \right\} [D(\lambda^*) - D(\lambda^k)]. \end{aligned}$$

266 From (29) and $\Delta_{k,l} > 0$, we have that $\lim_{k \rightarrow \infty} D(\lambda^*) - D(\lambda^k) = 0$. \square

267 **Theorem 2.4** implies finite convergence to an ϵ -optimum for $\epsilon > 0$. In other words,
 268 there exists $K \in \mathbb{Z}_+$ for any $\epsilon > 0$ such that $D(\lambda^*) - D(\lambda^k) < \epsilon$ for $k \geq K$. Also
 269 note that **Theorem 2.4** is valid for any $\Delta > 0$. The convergence analysis remains valid
 270 independent of the choice of the TR norm, as shown in the proof of **Lemma 2.3**. The
 271 convergence result is also independent of the bundle management strategy at serious
 272 steps. For instance, **Theorem 2.4** holds even if all the cuts are removed at every serious
 273 step. However, a suitable bundle management strategy may improve the efficiency of
 274 **Algorithm 1**.

275 **3. An Asynchronous Variant.** In this section we present an asynchronous
 276 implementation variant of the BTR algorithm described in **Algorithm 1**. To achieve
 277 asynchronicity, we only use a subset of scenario indices to update the master problem
 278 and the TR. Let $J^i \subseteq J$ be the subset of scenario indices such that bundle information
 279 $i \in \mathcal{B}^{k,l}$ is added to the model. We define the model function

$$280 \quad (30a) \quad \tilde{m}_{k,l}(\lambda) := \max_{j \in J} \sum \theta_j$$

$$281 \quad (30b) \quad \text{s.t. } \theta_j \leq D_j(\lambda^i) + (H_j x_j^i)^T (\lambda - \lambda^i), \forall i \in \mathcal{B}^{k,l}, j \in J^i,$$

283 where the cuts (30b) are generated only for a subset of scenarios J^i . Note that
 284 $\tilde{m}_{k,l}(\lambda) \geq D(\lambda)$ and $m_{k,l}(\lambda) \geq D(\lambda)$ for any given $\lambda \in \mathbb{R}^{N_{n_1}}$ and (k, l) .

285 At iteration (k, l) , the master problem of the asynchronous variant is given by

$$286 \quad (31) \quad \max_{\lambda \in \mathbb{R}^{N_{n_1}}} \{ \tilde{m}_{k,l}(\lambda) : \|\lambda - \lambda^k\| \leq \Delta_{k,l} \}.$$

288 While the master problem (31) is a natural extension of (6), it is not straightforward
 289 to guarantee convergence under this setting. In particular, existing algorithms assume
 290 full scenario synchronization before updating the TR (i.e., by checking (8), (16), and
 291 (18)), and terminating the algorithm (9). We now describe necessary modifications
 292 to ensure convergence.

293 We consider a set of processes that consist of a master process and multiple worker
 294 processes. The master process is responsible for solving the master problem (31) to
 295 find a new trial point $\lambda^{k,l}$ for the Lagrangian dual function $D(\cdot)$. We let Π denote
 296 the set of worker processes. Each worker process $\pi \in \Pi$ is responsible for solving La-
 297 grangian subproblems for a subset $J_\pi \subseteq J$ of scenarios to evaluate the trial point. We
 298 let $\Pi_{k,l} \subseteq \Pi$ denote a subset of idle worker processes ready for subproblem solutions
 299 at iteration (k, l) . We define $\underline{\Pi} \in (0, |\Pi|]$ as the minimum number of worker processes
 300 that are ready for subproblem solutions. We let $\Lambda_{k,l}$ denote a queue for trial points λ^q
 301 and status s^q , where q index iterations (k, l) , the status of evaluation of trial point λ^q
 302 is encoded in s_π^q for $\pi \in \Pi$, and $s^q := (s_\pi^q, \forall \pi \in \Pi)$. Each queue element is initialized
 303 by $s^q = \text{ready}$, and $s_\pi^q = \text{assigned}$ is encoded when λ^q is under evaluation at process
 304 π . Once λ^q is evaluated at process π , the status is set to $s_\pi^q = \text{evaluated}$. We let $\bar{\Lambda}$
 305 be the maximum number of elements in the queue $\Lambda_{k,l}$ at any iteration (k, l) .

306 We denote the predicted increase at iteration (k, l) by

$$307 \quad (32) \quad \tilde{v}_{k,l} := \tilde{m}_{k,l}(\lambda^{k,l}) - D(\lambda^k).$$

309 The serious steps and null steps are taken only when a trial point is evaluated by
 310 all worker processes. Otherwise, we update the model function $\tilde{m}_{k,l+1}(\cdot)$ by adding
 311 cuts (30b). In particular, if there exists

$$312 \quad (33) \quad \hat{\lambda}^{k,l} \in \arg \max \{ D(\lambda) : (\lambda, s) \in \Lambda_{k,l}, s_\pi = \text{evaluated} \forall \pi \in \Pi \}$$

313

314 and

$$315 \quad (34) \quad D(\hat{\lambda}^{k,l}) \geq D(\lambda^k) + \xi \tilde{v}_{k,l},$$

317 we update the TR center $\lambda^{k+1} \leftarrow \hat{\lambda}^{k,l}$. Otherwise, a null step is taken to update the
 318 model function $\tilde{m}_{k,l+1}(\cdot)$. We terminate the algorithm if

$$319 \quad (35) \quad \tilde{v}_{k,l} \leq \epsilon(1 + |D(\lambda^k)|).$$

321 We devise tests for adjusting the TR size $\Delta_{k,l}$ while still retaining convergence. For
 322 a small TR size, we can use the criterion (18) and update (19). In the asynchronous
 323 variant, not all scenario information is available at most null steps. To handle this
 324 case, we deem the TR to be too large if

$$325 \quad (36) \quad \tilde{\delta}^{k,l} := \frac{|J| \sum_{\pi \in \Pi_{k,l}} \sum_{j \in J_\pi} \delta_j^{k,l}}{\sum_{\pi \in \Pi_{k,l}} |J_\pi|} > \frac{\tilde{v}_{k,l}}{\min\{1, \|\lambda^{k,l} - \lambda^k\|\}}.$$

327 We define

$$328 \quad (37) \quad \tilde{\rho} := \min\{1, \|\lambda^{k,l} - \lambda^k\|\} \frac{\tilde{\delta}^{k,l}}{v_{k,l}},$$

330 and decrease the TR size as follows:

- 331 1. If $\tilde{\rho} > 0$, then $\tau^{k+1} \leftarrow \tau^k + 1$.
- 332 2. If $\tilde{\rho} > \bar{\rho}$ or $(\rho \in (0, \bar{\rho})$ and $\tau \geq \bar{\rho})$, then set $\tau^{k+1} \leftarrow 0$, and

$$333 \quad (38) \quad \Delta_{k,l+1} \leftarrow \max\left\{\frac{\Delta_{k,l}}{\min\{\rho, \bar{\rho}\}}, \Delta^{min}\right\}.$$

335 **3.1. Algorithmic Steps.** The steps of the asynchronous BTR algorithm are
 336 summarized in Algorithms 2 and 3. The steps in Algorithm 2 are taken in the master
 337 process and those in Algorithm 3 are taken in the worker processes. Algorithm 2 is
 338 initialized by setting parameters (line 1), collecting the initial bundle information (line
 339 2), and creating the model function $\tilde{m}_{0,0}$ (line 3). The main steps at each iteration
 340 (k, l) are described in the loop (lines 4–45). A new trial point $\lambda^{k,l}$ is obtained by
 341 solving the master (31) (line 5). If the termination criterion is satisfied by $\tilde{m}_{k,l}(\lambda^{k,l})$,
 342 the algorithm terminates (lines 6–8). Otherwise, the trial point is queued with the
 343 initial status $s_\pi^{k,l} \leftarrow \text{ready}$ for all $\pi \in \Pi$, if the queue $\Lambda_{k,l}$ is not full (lines 9–12). For
 344 each available worker process $\pi \in \Pi_{k,l}$, the master process chooses a trial point λ^q in
 345 queue $\Lambda_{k,l}$ for the Lagrangian subproblem solutions (line 15), if it exists, and sends it
 346 to the worker process (line 16). Here, we may consider two policies for choosing the
 347 trial points: first-in-first-out (FIFO) and last-in-first-out (LIFO). The convergence
 348 results in Section 3.2 are not affected by the choice of policy. If there exists no trial
 349 point to evaluate in the queue, the current trial point is sent to the worker process
 350 (line 19).

351 We note that the master process does not wait for all the worker processes to
 352 complete the evaluations. The master process receives the subproblem solutions from
 353 at least $\underline{\Pi}$ worker processes π and encodes $s_\pi^q \leftarrow \text{evaluated}$ for each λ^q and $\pi \in \Pi_{k,l}$
 354 if the trial point is from the queue (lines 23–27). The parameter $\underline{\Pi}$ controls the
 355 frequency of each master problem solve. It is possible that $|\Pi_{k,l}| > \underline{\Pi}$ if the previous

iteration takes time for completing the evaluations in a large enough number of worker processes. For the trial points evaluated by all worker processes, the master process chooses the best trial point for the serious step test (8) (line 30). If there exists a trial point evaluated by all worker processes, the TR size may be updated (lines 33 and 38). If the serious step test is satisfied, the master process updates the TR center (line 35). If there exists no trial point evaluated by all worker processes, the model function $\tilde{m}_{k,l+1}$ is updated by adding cuts (line 42). Each worker process π repeats the steps in lines 1 – 5 until receiving the termination signal from the master process. In the loop, the worker process receives a new trial point λ from the master process (line 2), evaluates it (line 3), and sends the information $D_j(\lambda), x_j$ to the master process (line 4).

3.2. Convergence Analysis. We analyze the convergence of Algorithm 2. We note that for $\underline{\Pi} = |\Pi|$, Algorithm 2 is equivalent to Algorithm 1. Therefore, we assume that $\underline{\Pi} \in (0, |\Pi|)$ and we assume that $\epsilon = 0$ throughout this section. At every major iteration $k \geq 0$, we have that $\tilde{m}_{k,l+1}(\lambda) \leq \tilde{m}_{k,l}(\lambda)$ for all $l \geq 0$.

We first show that Algorithm 2 takes only a finite number of null steps.

LEMMA 3.1. *Suppose that Algorithm 2 takes a null step at iteration $(k, l_{(1)})$. There exist a finite number $i > 1$ such that the algorithm either makes a serious step or terminates at iteration $(k, l_{(i)})$.*

Proof. Suppose for contradiction that the algorithm takes null steps at iteration $(k, l_{(i)})$ for all $i \geq 1$. Then, we have for $i \geq 1$

$$(39) \quad \tilde{m}_{k,l_{(i)}}(\lambda^{k,l_{(i)}}) - D(\lambda^k) > \epsilon(1 + |D(\lambda^k)|) = 0$$

and

$$(40) \quad D(\hat{\lambda}^{k,l_{(i)}}) < D(\lambda^k) + \xi [\tilde{m}_{k,l_{(i)}}(\lambda^{k,l_{(i)}}) - D(\lambda^k)].$$

Let (k_i, l_i) be the iteration at which $\hat{\lambda}^{k,l_{(i)}}$ was obtained by the master (31) so that $\lambda^{k_i,l_i} = \hat{\lambda}^{k,l_{(i)}}$ for $i \geq 1$. For the case that $k_i < k$, the algorithm can find $\hat{\lambda}^{k,l_{(i)}}$ such that $k_I = k$ and $l_I = l_{(i)}$ for $i < I < \infty$. Therefore, we need only to consider the case that $k_i = k$ and $l_i \leq l_{(i)}$. In such a case, there exists some $j \in J$ such that the linear inequalities generated at λ^{k_i,l_i} ,

$$(41) \quad \theta_j \leq D(\lambda^{k_i,l_i}) + (H_j x_j^{k_i,l_i})^T (\lambda - \lambda^{k_i,l_i}),$$

are violated at $(\theta^{k_i,l_i}, \lambda^{k_i,l_i})$, because

$$\begin{aligned} \sum_{j \in J} \theta_j^{k_i,l_i} - D(\lambda^{k_i,l_i}) &= \tilde{m}_{k_i,l_i}(\lambda^{k_i,l_i}) - D(\lambda^{k_i,l_i}) \\ &> \tilde{m}_{k_i,l_i}(\lambda^{k_i,l_i}) - D(\lambda^k) - \xi [\tilde{m}_{k,l_{(i)}}(\lambda^{k,l_{(i)}}) - D(\lambda^k)] \\ &\geq \tilde{m}_{k_i,l_i}(\lambda^{k,l_{(i)}}) - D(\lambda^k) - \xi [\tilde{m}_{k,l_{(i)}}(\lambda^{k,l_{(i)}}) - D(\lambda^k)] \\ &\geq \tilde{m}_{k,l_{(i)}}(\lambda^{k,l_{(i)}}) - D(\lambda^k) - \xi [\tilde{m}_{k,l_{(i)}}(\lambda^{k,l_{(i)}}) - D(\lambda^k)] \\ &> 0. \end{aligned}$$

Here, the second inequality holds because λ^{k_i,l_i} is a maximizer of \tilde{m}_{k_i,l_i} and the third inequality holds because $\tilde{m}_{k,l}$ is nonincreasing within a given major iteration k . From Lemma 2.1, the algorithm violates either (39) or (40). \square

Algorithm 2 Asynchronous Bundle-Trust-Region Algorithm

1: Initialize $\lambda^0 \in \mathbb{R}^{Nn_1}$, $\Delta_{0,0} \in (0, \Delta^{max}]$, $\xi \in (0, 0.5)$, $\epsilon \geq 0$, $\bar{\Lambda} > 0$, $\underline{\Pi} \in (0, |\Pi|]$, $\Lambda_{0,0} \leftarrow \emptyset$, $\Pi_{0,0} \leftarrow \Pi$, $\mathcal{B}^{0,0} \leftarrow \{0\}$, $k \leftarrow 0$, and $l \leftarrow 0$.
 2: Solve the Lagrangian subproblem (4) to find $D_j(\lambda^0)$ and x_j^0 for all $j \in J$.
 3: Initialize the model function $\tilde{m}_{0,0}$.
 4: **loop**
 5: Solve the master (31) to find $\lambda^{k,l}$.
 6: **if** $\tilde{m}_{k,l}(\lambda^{k,l}) - D(\lambda^k) \leq \epsilon(1 + |D(\lambda^k)|)$ **then** ▷ Termination test
 7: Stop
 8: **end if**
 9: **if** $|\Lambda_{k,l}| < \bar{\Lambda}$ **then** ▷ Queue new dual variable
 10: $s_\pi^{k,l} \leftarrow \text{ready}$ for all $\pi \in \Pi$.
 11: $\Lambda_{k,l} \leftarrow \Lambda_{k,l} \cup \{(\lambda^{k,l}, s^{k,l})\}$.
 12: **end if**
 13: **for** $\pi \in \Pi_{k,l}$ **do** ▷ Send dual variables to processes
 14: **if** $\exists(\lambda^q, s^q) \in \Lambda_{k,l}$ such that $s_\pi^q = \text{ready}$ **then**
 15: Choose an element $(\lambda^q, s^q) \in \Lambda_{k,l}$ such that $s_\pi^q = \text{ready}$.
 16: SEND λ^q to process π .
 17: $s_\pi^q \leftarrow \text{assigned}$.
 18: **else**
 19: SEND $\lambda^{k,l}$ to process π .
 20: **end if**
 21: $\Pi_{k,l} \leftarrow \Pi_{k,l} \setminus \{\pi\}$.
 22: **end for**
 23: **repeat** ▷ Receive subproblem solutions from processes
 24: RECEIVE $D_j(\lambda^q)$ and x_j^q for $j \in J_\pi$ from any process $\pi \in \Pi$.
 25: $s_\pi^q \leftarrow \text{evaluated}$ if $(\lambda^q, \cdot) \in \Lambda_{k,l}$.
 26: $\Pi_{k,l} \leftarrow \Pi_{k,l} \cup \{\pi\}$.
 27: **until** $|\Pi_{k,l}| \geq \underline{\Pi}$
 28: **serious** $\leftarrow \text{false}$.
 29: **if** $\exists(\lambda, s) \in \Lambda_{k,l}$ such that $s_\pi = \text{evaluated} \forall \pi \in \Pi$ **then**
 30: $\hat{\lambda}^{k,l} \leftarrow \arg \max\{D(\lambda) : (\lambda, s) \in \Lambda_{k,l}, s_\pi = \text{evaluated} \forall \pi \in \Pi\}$
 31: $\Lambda_{k,l} \leftarrow \Lambda_{k,l} \setminus \{(\hat{\lambda}^{k,l}, \text{evaluated})\}$
 32: **if** $D(\hat{\lambda}^{k,l}) \geq D(\lambda^k) + \xi[\tilde{m}_{k,l}(\lambda^{k,l}) - D(\lambda^k)]$ **then** ▷ Serious step
 33: Choose $\Delta_{k+1,0} \in [\Delta_{k,l}, \Delta^{max}]$.
 34: Choose $\Lambda_{k+1,0} \subseteq \Lambda_{k,l}$.
 35: Set $\lambda^{k+1} \leftarrow \hat{\lambda}^{k,l}$, $\tilde{m}_{k+1,0} \leftarrow \tilde{m}_{k,l}$, $\Pi_{k+1,0} \leftarrow \Pi_{k,l}$, $k \leftarrow k + 1$ and $l \leftarrow 0$.
 36: **serious** $\leftarrow \text{true}$.
 37: **else** ▷ Null step
 38: Choose $\Delta_{k,l+1} \in (0, \Delta_{k,l}]$.
 39: **end if**
 40: **end if**
 41: **if** **serious** = **false** **then** ▷ Model update
 42: Update the model function $\tilde{m}_{k,l+1}$ by adding cuts (30b).
 43: Set $\Pi_{k,l+1} \leftarrow \Pi_{k,l}$, $\Lambda_{k,l+1} \leftarrow \Lambda_{k,l}$, $l \leftarrow l + 1$.
 44: **end if**
 45: **end loop**

Algorithm 3 Asynchronous Bundle-Trust-Region Algorithm - Worker (π)

- 1: **repeat**
 - 2: RECEIVE new trial point λ from the master process.
 - 3: Solve the Lagrangian subproblem (4) to find $D_j(\lambda)$ and x_j for all $j \in J_\pi$.
 - 4: SEND $D_j(\lambda), x_j$ for $j \in J_\pi$ to the master process.
 - 5: **until** the master process terminates.
-

399 Analogous to [Lemma 2.3](#), we derive a lower bound for the predicted increase in
400 the lower bound.

401 **LEMMA 3.2.** *For $k \geq 0$ and $l \geq 0$, we have that*

$$402 \quad (42) \quad \tilde{m}_{k,l}(\lambda^{k,l}) - \tilde{m}_{k,l}(\lambda^k) \geq \min \left\{ \frac{\Delta_{k,l}}{\|\lambda^* - \lambda^k\|}, 1 \right\} [D(\lambda^*) - D(\lambda^k)].$$

404 *Proof.* The proof follows the steps in the proof of [Lemma 2.3](#). □

405 We now show that [Algorithm 2](#) finds the Lagrangian dual bound in the limit.

406 **THEOREM 3.3.** *Algorithm 2 delivers a sequence of dual iterates $\{\lambda^k\}$ satisfying*
407 $\lim_{k \rightarrow \infty} D(\lambda^k) \rightarrow D(\lambda^*)$.

408 *Proof.* Let l_k be such that $\hat{\lambda}^{k,l_k}$ takes a serious step and thus $\lambda^{k+1} = \hat{\lambda}^{k,l_k}$. Since

$$409 \quad (43) \quad \tilde{m}_{k,l_k}(\lambda^{k,l_k}) - D(\lambda^k) \geq \epsilon(1 + |D(\lambda^k)|) > 0$$

411 and

$$412 \quad (44) \quad D(\hat{\lambda}^{k,l_k}) - D(\lambda^k) \geq \xi[\tilde{m}_{k,l_k}(\lambda^{k,l_k}) - D(\lambda^k)] > 0,$$

414 we have $D(\lambda^{k+1}) - D(\lambda^k) > 0$. Since $\{D(\lambda^k)\}$ is an increasing sequence that is
415 bounded above by z_{LD} , we have that

$$416 \quad (45) \quad \lim_{k \rightarrow \infty} D(\lambda^{k+1}) - D(\lambda^k) = 0.$$

418 Moreover, by [Lemma 3.2](#) we have

$$419 \quad \tilde{m}_{k,l_k}(\hat{\lambda}^{k,l_k}) - \tilde{m}_{k,l_k}(\lambda^k) = D(\lambda^{k+1}) - D(\lambda^k) \\ 420 \quad \geq \min \left\{ \frac{\Delta_{k,l_k}}{\|\lambda^* - \lambda^k\|}, 1 \right\} [D(\lambda^*) - D(\lambda^k)].$$

422 From (45) and $\Delta_{k,l} > 0$, we have $\lim_{k \rightarrow \infty} D(\lambda^*) - D(\lambda^k) = 0$. □

423 **3.3. Dynamic Subproblem Allocation.** [Algorithm 2](#) runs for a fixed (static)
424 set J_π for each $\pi \in \Pi$, where each worker process is allocated to certain scenario
425 subproblems for all iterations. The allocation of worker processes is advantageous
426 because each MIP subproblem solver can take advantage of the warm-start feature
427 given in off-the-shelf MIP solvers. However, the static allocation of subproblems
428 can cause parallel inefficiency, because one of the worker processes might present a
429 computational bottleneck for evaluating the trial points in queue.

430 The asynchronous computation in [Algorithm 2](#) can be varied by dynamically allo-
431 cating the subproblems to the worker processes. With dynamic allocation, the worker
432 processes are not dedicated to certain scenario subproblems and can possibly solve

433 different scenario subproblems at each iteration. An asynchronous algorithm with dy-
 434 namic allocation can be obtained by performing minor modifications to [Algorithm 2](#).
 435 In particular, we only need to keep track of the status of the evaluation of the trial
 436 points in the queue for each scenario subproblem, as compared with that for each
 437 worker process in [Algorithm 2](#). Similar to the definition of s_π^q used in [Algorithm 2](#),
 438 we define \tilde{s}_j^q as the status of evaluation of trial point λ^q for scenario index $j \in J$, and
 439 $\tilde{s}^q := (\tilde{s}_j^q, \forall j \in J)$. The modified steps are summarized in [Algorithm 4](#).

440 The initial queue status $\tilde{s}_j^{k,l}$ is set for each scenario $j \in J$ in line 10. We choose
 441 a trial point that is not evaluated for some scenarios and send it to a worker process
 442 (lines 14–17). For the choice of trial points, we are interested in FIFO and LIFO poli-
 443 cies, as mentioned in [Subsection 3.1](#). The rest of the algorithm was modified in order
 444 to apply the modified notation \tilde{s} for updating and checking the status of evaluating
 445 trial points for each scenario (lines 25, 29–30). We note that the convergence analysis
 446 given in [Subsection 3.2](#) holds for [Algorithm 4](#).

447 **4. Numerical Experiments.** In this section, we present numerical experiments
 448 for the synchronous and asynchronous BTR algorithms.

449 **4.1. Implementation.** We have implemented the proposed algorithms in the
 450 open-source and parallel software package DSP [14]. The master problem and MIP
 451 subproblems were solved using CPLEX 12.7. The master problem was solved using a
 452 barrier method with 16 parallel cores, whereas the MIP subproblems were solved with
 453 default parameter settings. Moreover, the subproblems were solved in parallel by us-
 454 ing MPI, and each process uses a single core. The subproblems are distributed to the
 455 processes in a round-and-robin fashion (except for the dynamic allocation strategy in
 456 [Subsection 4.5](#)). The processes create a CPLEX solver environment for each subprob-
 457 lem in order to take advantage of the CPLEX warm-starting feature, as well as to avoid
 458 data loading time. All computations were performed on the *Blues* cluster — a 630-
 459 node computing cluster at Argonne National Laboratory. For both the synchronous
 460 and asynchronous BTR algorithms, we set $\lambda^0 = 0$. We also set the parameters $\bar{\rho} = 3$,
 461 $\underline{\rho} = 4$, $\underline{\Delta} = 10^{-2}$, $\bar{\Delta} = 10^4$, $\Delta_{0,0} = 10^2$, $\xi = 10^{-4}$, $\epsilon = 10^{-5}$. We considered different
 462 values for parameters $\bar{\Lambda}$, $\underline{\Pi}$ of [Algorithm 2](#) to evaluate performance, as discussed in
 463 [Subsection 4.4](#).

464 **4.2. Problem Instances.** We use a day-ahead stochastic unit commitment
 465 (SUC) problem with data representing a test system for the California independent
 466 system operator (CAISO) interconnected with the Western Electricity Coordinating
 467 Council, as discussed in [21, 13, 11]. The SUC problem instances embed difficult MIP
 468 scenario subproblems that induce strong load imbalances. The test system consists of
 469 225 buses, 375 transmission lines, 130 generators, 40 loads, 5 import points, 5 wind
 470 farms, and 11 non-wind renewable generators. Of the 130 generators, 90 generators
 471 are *fast generators* capable of starting in response to demand. The other 40 generators
 472 are *slow generators*. In the SUC problem instances, we consider a 24-hour time horizon
 473 with hourly intervals. The amount of power at import points and renewable genera-
 474 tors and loads are categorized in eight day types. Each day type is a combination of
 475 elements in two sets $\{Spring, Summer, Fall, Winter\}$ and $\{Weekdays, Weekends\}$.
 476 The load is calculated by the amount of total system load subtracted by total import
 477 power and renewable power, excluding wind power (see [Figure 1](#)). In addition to the
 478 eight day types, we use 160 scenarios of wind power generation at each wind farm and
 479 for each season (see [Figure 2](#)).

480 We use \mathcal{N} , \mathcal{L} , \mathcal{T} , and \mathcal{S} to represent sets of buses, transmission lines, time periods,

Algorithm 4 Asynchronous Bundle-Trust-Region Algorithm with Dynamic Subproblem Allocation

```

1: Initialize  $\lambda^0 \in \mathbb{R}^{Nn_1}, \Delta_{0,0} \in (0, \Delta^{max}], \xi \in (0, 0.5), \epsilon \geq 0, \bar{\Lambda} > 0, \underline{\Pi} \in$ 
    $(0, |\underline{\Pi}|], \Lambda_{0,0} \leftarrow \emptyset, \Pi_{0,0} \leftarrow \underline{\Pi}, \mathcal{B}^{0,0} \leftarrow \{0\}, k \leftarrow 0,$  and  $l \leftarrow 0.$ 
2: Solve the Lagrangian subproblem (4) to find  $D_j(\lambda^0)$  and  $x_j^0$  for all  $j \in J.$ 
3: Initialize the model function  $\tilde{m}_{0,0}.$ 
4: loop
5:   Solve the master (31) to find  $\lambda^{k,l}.$ 
6:   if  $\tilde{m}_{k,l}(\lambda^{k,l}) - D(\lambda^k) \leq \epsilon(1 + |D(\lambda^k)|)$  then ▷ Termination test
7:     Stop
8:   end if
9:   if  $|\Lambda_{k,l}| < \bar{\Lambda}$  then ▷ Queue new dual variable
10:      $\tilde{s}_j^{k,l} \leftarrow \text{ready}$  for all  $j \in J.$ 
11:      $\Lambda_{k,l} \leftarrow \Lambda_{k,l} \cup \{(\lambda^{k,l}, \tilde{s}^{k,l})\}.$ 
12:   end if
13:   for  $\pi \in \Pi_{k,l}$  do ▷ Dynamic allocation of dual variables to processes
14:     if  $\exists(\lambda^q, s^q) \in \Lambda_{k,l}$  such that  $\tilde{s}_j^q = \text{ready}$  for any  $j \in J$  then
15:       Choose an element  $(\lambda^q, \tilde{s}^q) \in \Lambda_{k,l}$  such that  $\tilde{s}_j^q = \text{ready}$  for some  $j \in J.$ 
16:       SEND  $\lambda^q$  to process  $\pi.$ 
17:        $\tilde{s}_j^q \leftarrow \text{assigned}$  and  $J_\pi \leftarrow \{j\}.$ 
18:     else
19:       SEND  $\lambda^{k,l}$  to process  $\pi.$ 
20:     end if
21:      $\Pi_{k,l} \leftarrow \Pi_{k,l} \setminus \{\pi\}.$ 
22:   end for
23:   repeat ▷ Receive subproblem solutions from processes
24:     RECEIVE  $D_j(\lambda^q)$  and  $x_j^q$  from any process  $\pi \in \Pi.$ 
25:      $\tilde{s}_j^q \leftarrow \text{evaluated}$  if  $(\lambda^q, \cdot) \in \Lambda_{k,l}.$ 
26:      $\Pi_{k,l} \leftarrow \Pi_{k,l} \cup \{\pi\}.$ 
27:   until  $|\Pi_{k,l}| \geq \underline{\Pi}$ 
28:   serious  $\leftarrow \text{false}.$ 
29:   if  $\exists(\lambda, \tilde{s}) \in \Lambda_{k,l}$  such that  $\tilde{s}_j = \text{evaluated} \forall j \in J$  then
30:      $\hat{\lambda}^{k,l} \leftarrow \arg \max\{D(\lambda) : (\lambda, \tilde{s}) \in \Lambda_{k,l}, \tilde{s}_j = \text{evaluated} \forall j \in J\}$ 
31:      $\Lambda_{k,l} \leftarrow \Lambda_{k,l} \setminus \{(\hat{\lambda}^{k,l}, \text{evaluated})\}$ 
32:     if  $D(\hat{\lambda}^{k,l}) \geq D(\lambda^k) + \xi[\tilde{m}_{k,l}(\lambda^{k,l}) - D(\lambda^k)]$  then ▷ Serious step
33:       Choose  $\Delta_{k+1,0} \in [\Delta_{k,l}, \Delta^{max}].$ 
34:       Choose  $\Lambda_{k+1,0} \subseteq \Lambda_{k,l}.$ 
35:       Set  $\lambda^{k+1} \leftarrow \hat{\lambda}^{k,l}, \tilde{m}_{k+1,0} \leftarrow \tilde{m}_{k,l}, \Pi_{k+1,0} \leftarrow \Pi_{k,l}, k \leftarrow k + 1$  and  $l \leftarrow 0.$ 
36:       serious  $\leftarrow \text{true}.$ 
37:     else ▷ Null step
38:       Choose  $\Delta_{k,l+1} \in (0, \Delta_{k,l}].$ 
39:     end if
40:   end if
41:   if serious = false then ▷ Model update
42:     Update the model function  $\tilde{m}_{k,l+1}$  by adding cuts (30b).
43:     Set  $\Pi_{k,l+1} \leftarrow \Pi_{k,l}, \Lambda_{k,l+1} \leftarrow \Lambda_{k,l}, l \leftarrow l + 1.$ 
44:   end if
45: end loop

```




Fig. 1: Net load of the test system for the 24-hour time horizon for each day type. WD and WE in the legend stand for weekdays and weekends, respectively.




Fig. 2: Total amount of wind power for each season. Scenarios are shown in grey and the mean is in blue.

481 and scenarios, respectively. In addition, \mathcal{L}_n^+ and $\mathcal{L}_n^- \subset \mathcal{L}$ represent the set of trans-
 482 mission lines to and from bus $n \in \mathcal{N}$, respectively. We use $\mathcal{G}_n, \mathcal{D}_n, \mathcal{I}_n, \mathcal{R}_n$, and \mathcal{W}_n
 483 to denote sets of generators, loads, import points, renewable generators, and wind farms
 484 at bus $n \in \mathcal{N}$, respectively. We also define $\mathcal{G} := \sum_{n \in \mathcal{N}} \mathcal{G}_n$ and denote by \mathcal{G}_S the
 485 set of slow generators that should be scheduled a day ahead. Let u_{gts}, v_{gts} , and p_{gts}
 486 be the decision variables for unit commitment, generator start, and power generation
 487 amount, respectively, for $g \in \mathcal{G}, t \in \mathcal{T}, s \in \mathcal{S}$. Let f_{lts} be the decision variables for
 488 power flow on transmission line $l \in \mathcal{L}$, and let θ_{nts} be the decision variables for phase
 489 angles at bus $n \in \mathcal{N}$. Let $d_{jts}, m_{its}, r_{its}$, and w_{its} be the slack variables representing
 490 load shedding, import spillage, renewable generation spillage, and wind generation
 491 spillage, respectively.

492 The formulation of the SUC problem is given by the following two-stage stochastic
 493 mixed-binary program,

$$494 \quad (46a) \quad \min \sum_{t \in \mathcal{T}} \sum_{s \in \mathcal{S}} \frac{1}{|\mathcal{S}|} \left[\sum_{g \in \mathcal{G}} (K_g u_{gts} + S_g v_{gts} + C_g p_{gts}) + \sum_{j \in \mathcal{D}} V d_{jts} \right]$$

$$495 \quad (46b) \quad \text{s.t. } u_{gts} = u_{gt, s-1} \quad \forall g \in \mathcal{G}_S, t \in \mathcal{T}, s \in \mathcal{S},$$

$$496 \quad (46c) \quad v_{gts} = v_{gt, s-1} \quad \forall g \in \mathcal{G}_S, t \in \mathcal{T}, s \in \mathcal{S},$$

$$497 \quad (46d) \quad \sum_{q=t-UT_g+1}^t v_{gqs} \leq u_{gts} \quad \forall g \in \mathcal{G}, t \in \{UT_g, \dots, T\}, s \in \mathcal{S},$$

$$498 \quad (46e) \quad \sum_{q=t+1}^{t+DT_g} v_{gqs} \leq u_{gts} \quad \forall g \in \mathcal{G}, t \in \{1, \dots, T - DT_g\}, s \in \mathcal{S},$$

$$499 \quad (46f) \quad v_{gts} \geq u_{gts} - u_{g, t-1, s} \quad \forall g \in \mathcal{G}, t \in \mathcal{T}, s \in \mathcal{S},$$

$$500 \quad \sum_{l \in \mathcal{L}_n^+} f_{lts} - \sum_{l \in \mathcal{L}_n^-} f_{lts} + \sum_{g \in \mathcal{G}_n} p_{gts} = \sum_{j \in \mathcal{D}_n} (D_{jt} - d_{jts})$$

$$501 \quad - \sum_{i \in \mathcal{I}_n} (M_{it} - m_{its}) - \sum_{i \in \mathcal{R}_n} (R_{it} - r_{its})$$

$$502 \quad (46g) \quad - \sum_{i \in \mathcal{W}_n} (W_{its} - w_{its}) \quad \forall n \in \mathcal{N}, t \in \mathcal{T}, s \in \mathcal{S},$$

$$503 \quad (46h) \quad f_{lts} = B_l(\theta_{mts} - \theta_{nts}) \quad \forall l = (m, n) \in \mathcal{L}, t \in \mathcal{T}, s \in \mathcal{S},$$

$$504 \quad (46i) \quad u_{gts} \in \{0, 1\}, v_{gts} \in [0, 1], \quad \forall g \in \mathcal{G}, t \in \mathcal{T}, s \in \mathcal{S},$$

$$505 \quad (46j) \quad P_g^{\min} u_{gts} \leq p_{gts} \leq P_g^{\max} u_{gts}, \quad \forall g \in \mathcal{G}, t \in \mathcal{T}, s \in \mathcal{S},$$

$$506 \quad (46k) \quad -F^{\max} \leq f_{lts} \leq F^{\max}, \quad \forall l \in \mathcal{L}, t \in \mathcal{T}, s \in \mathcal{S},$$

$$507 \quad (46l) \quad -360 \leq \theta_{nts} \leq 360, \quad \forall n \in \mathcal{N}, t \in \mathcal{T}, s \in \mathcal{S},$$

$$508 \quad (46m) \quad 0 \leq d_{jts} \leq D_{jt}, \quad \forall j \in \mathcal{D}_n, n \in \mathcal{N}, t \in \mathcal{T}, s \in \mathcal{S},$$

$$509 \quad (46n) \quad 0 \leq m_{its} \leq M_{it}, \quad \forall i \in \mathcal{I}_n, n \in \mathcal{N}, t \in \mathcal{T}, s \in \mathcal{S},$$

$$510 \quad (46o) \quad 0 \leq r_{its} \leq R_{it}, \quad \forall i \in \mathcal{R}_n, n \in \mathcal{N}, t \in \mathcal{T}, s \in \mathcal{S},$$

$$511 \quad (46p) \quad 0 \leq w_{its} \leq W_{its}, \quad \forall i \in \mathcal{W}_n, n \in \mathcal{N}, t \in \mathcal{T}, s \in \mathcal{S},$$

513 where equations (46b) and (46c) are the nonanticipativity constraints on the first-
 514 stage variables u_{gts} and v_{gts} for $g \in \mathcal{G}_S, t \in \mathcal{T}, s \in \mathcal{S}$ (i.e., commitment decisions for
 515 slow generators). The objective function (46a) minimizes the expected total operat-
 516 ing cost that sums commitment cost K_g , generator startup cost S_g , generation cost

517 C_g , and load shedding cost V (set to \$5,000/MWh). Equations (46d) and (46e) are,
 518 respectively, the minimum uptime UT_g and downtime DT_g constraints for each gen-
 519 erator $g \in \mathcal{G}$. Constraint (46f) imposes the unit commitment logic. Constraint (46g)
 520 represents the power balance equation with load D_{jt} , import power M_{it} , renewable
 521 generation R_{it} , and wind generation W_{its} . Constraint (46h) represents the direct
 522 current power flow equation with line susceptance B_l . Constraints (46j) and (46k)
 523 impose the minimum P_g^{min} and maximum P_g^{max} generation capacity and transmission
 524 line capacity F^{max} , respectively.

525 Using the net load data and the wind generation scenarios, we created 80 SUC
 526 problem instances. Each instance uses 16 scenarios for wind power generation. The
 527 scenario subproblems are distributed in a round-and-robin fashion and we also explore
 528 the dynamic allocation described in Subsection 4.5. For each problem instance, the
 529 first stage has 1,920 variables and 1,960 constraints, and the second stage has 21,144
 530 variables and 21,930 constraints. Therefore, each SUC instance has 340,224 variables
 531 and 352,840 constraints.

532 **4.3. Synchronous Computing and Load Balance.** We first present compu-
 533 tational times and load balancing for the synchronous BTR algorithm, to highlight
 534 the impact of load imbalancing on efficiency. We solve the 80 problem instances by
 535 using DSP with 32 cores on the *Blues* cluster. The master process uses 16 cores, and
 536 worker processes use the other 16 cores for solving MIP subproblems in parallel. We
 537 define the set of problem instances as \mathcal{P} . The load balance is quantified by using
 538 the percent imbalance metric, as defined in [22]. Specifically, we define the percent
 539 imbalance metric of problem instance $p \in \mathcal{P}$ and for each iteration k as

$$(47) \quad \nu_{pk} := \left(\frac{t_{pk}^{max}}{\bar{t}_{pk}} - 1 \right) \times 100\%,$$

542 where t_{pk}^{max} and \bar{t}_{pk} are the maximum and mean subproblem solution times, respec-
 543 tively, for problem instance p over all worker processes at iteration k . We also define
 544 $\bar{\nu}_p := \max_k \nu_{pk}$ and $\underline{\nu}_p := \min_k \nu_{pk}$. We denote by ν_p^m the mean percent imbalance
 545 metric over all iterations. The computational results are reported in Table 1, where t
 546 represents the total solution time (in seconds) and t^{LB} denotes the total time spent
 547 in the computation of the lower bound.

Table 1: Computational results and percent imbalance metric from the synchronous BTR algorithm.

Instance	Iter	z_{LD}	t	t^{LB}	$\bar{\nu}_p$	$\underline{\nu}_p$	ν_p^m
SpringWD0	125	1602068	2439	2369	86%	17%	51%
SpringWD1	141	1799653	3026	2937	98%	20%	42%
SpringWD2	108	1787644	2046	1989	56%	15%	29%
SpringWD3	130	1726987	2567	2489	82%	18%	41%
SpringWD4	165	1898284	3182	3061	74%	18%	35%
SpringWD5	132	1999293	2767	2691	57%	14%	31%
SpringWD6	115	1615039	2170	2109	108%	22%	47%
SpringWD7	196	1607014	5209	4963	138%	19%	73%
SpringWD8	179	1725101	4134	3932	91%	22%	47%
SpringWD9	126	1782099	2433	2357	94%	21%	42%

Continued on next page

Table 1 – *Continued from the previous page*

Instance	Iter	z_{LD}	t	t^{LB}	$\bar{\nu}_p$	$\underline{\nu}_p$	ν_p^m
SpringWE0	140	1044900	2397	2337	77%	13%	36%
SpringWE1	121	1114237	1826	1772	73%	10%	32%
SpringWE2	123	1110895	2211	2161	96%	21%	43%
SpringWE3	132	1084464	2410	2345	93%	18%	51%
SpringWE4	111	1199183	1653	1608	77%	12%	30%
SpringWE5	118	1259176	1732	1681	62%	15%	31%
SpringWE6	176	974356	2960	2700	98%	18%	42%
SpringWE7	154	1007195	2682	2598	94%	18%	47%
SpringWE8	138	1090599	2252	2181	76%	10%	39%
SpringWE9	121	1133203	1899	1760	96%	17%	37%
SummerWD0	189	4710799	5086	4874	110%	18%	50%
SummerWD1	149	4753046	4269	4225	115%	20%	48%
SummerWD2	112	4800483	2798	2771	121%	19%	52%
SummerWD3	208	4693318	5948	5532	97%	19%	50%
SummerWD4	149	4743635	3189	3019	60%	12%	34%
SummerWD5	288	4549022	10782	6942	164%	0%	82%
SummerWD6	175	4698568	5719	5497	174%	19%	85%
SummerWD7	134	4644613	3584	3449	103%	24%	51%
SummerWD8	115	4827235	2516	2487	64%	17%	34%
SummerWD9	276	4678698	8794	8181	123%	23%	60%
SummerWE0	160	3159044	3176	3010	61%	15%	34%
SummerWE1	85	3200891	1408	1390	75%	16%	31%
SummerWE2	80	3237591	1406	1390	77%	14%	32%
SummerWE3	129	3162672	2182	2078	67%	12%	28%
SummerWE4	91	3208814	1627	1572	51%	12%	29%
SummerWE5	163	3057287	2992	2838	111%	17%	42%
SummerWE6	151	3141373	2621	2477	61%	14%	30%
SummerWE7	111	3108884	2027	1955	71%	15%	30%
SummerWE8	90	3248660	1719	1662	66%	15%	34%
SummerWE9	206	3161301	3876	3503	63%	13%	32%
FallWD0	165	2444953	5421	5288	75%	29%	52%
FallWD1	150	2555819	5897	5768	129%	19%	63%
FallWD2	144	2495699	4225	4120	101%	25%	49%
FallWD3	265	2470042	10244	9543	123%	23%	61%
FallWD4	175	2562315	5528	5338	95%	23%	48%
FallWD5	184	2563568	5260	5059	92%	23%	48%
FallWD6	159	2481434	4468	4340	79%	15%	46%
FallWD7	154	2624963	4052	3916	113%	18%	54%
FallWD8	164	2525282	4746	4582	84%	15%	43%
FallWD9	168	2406007	4498	4331	90%	23%	48%
FallWE0	121	1542897	1925	1860	75%	16%	40%
FallWE1	108	1622543	1705	1649	76%	16%	37%
FallWE2	149	1550416	3040	2918	75%	22%	44%
FallWE3	134	1541075	2436	2357	100%	22%	46%
FallWE4	107	1627667	1849	1796	79%	16%	39%
FallWE5	115	1631388	1903	1751	63%	16%	35%

Continued on next page

Table 1 – *Continued from the previous page*

Instance	Iter	z_{LD}	t	t^{LB}	$\bar{\nu}_p$	$\underline{\nu}_p$	ν_p^m
FallWE6	112	1567312	1785	1727	71%	14%	38%
FallWE7	97	1687574	1835	1788	92%	21%	44%
FallWE8	101	1589716	1791	1741	102%	13%	43%
FallWE9	113	1503765	1775	1718	75%	18%	35%
WinterWD0	146	1869441	2994	2899	71%	18%	39%
WinterWD1	147	1952864	3243	3141	130%	25%	58%
WinterWD2	110	1674367	1925	1866	105%	20%	47%
WinterWD3	136	1786630	2923	2839	125%	29%	58%
WinterWD4	117	1735955	1992	1926	80%	17%	38%
WinterWD5	142	1932486	2937	2844	73%	21%	38%
WinterWD6	137	2036682	2941	2841	74%	16%	42%
WinterWD7	143	1994766	3170	3075	70%	19%	40%
WinterWD8	138	1738391	2656	2566	97%	25%	47%
WinterWD9	136	1733211	2452	2364	103%	19%	41%
WinterWE0	130	1293480	2108	2038	83%	16%	38%
WinterWE1	147	1341952	2556	2461	87%	16%	36%
WinterWE2	129	1166758	1985	1926	64%	16%	36%
WinterWE3	197	1238547	3505	3302	119%	23%	60%
WinterWE4	119	1209598	1821	1768	72%	16%	39%
WinterWE5	156	1336934	2693	2582	94%	21%	44%
WinterWE6	122	1403328	1843	1778	78%	18%	38%
WinterWE7	141	1375098	2341	2255	101%	18%	38%
WinterWE8	110	1187088	1770	1724	75%	23%	41%
WinterWE9	197	1207103	5494	5249	233%	22%	76%

548 For the 80 problem instances, the total solution time ranges from 1,406 to 10,782
 549 seconds. The average solution time is 3,193 seconds, of which 3118 seconds were
 550 spent on the lower bounding problem. The average lower bounding time per iteration
 551 is 20 seconds. Figure 3 summarizes the distribution of the average percent imbalance
 552 metrics, where the y-axis presents the number of problem instances such that the
 553 percent imbalance metric is larger than or equal to the x-axis value. The average
 554 percent imbalance metrics range from 27% to 84%. The average percent imbalance is
 555 larger than 50% for the 18 problem instances.

556 **4.4. Asynchronous Computing and Performance Profiles.** We analyze
 557 the computational performance of the asynchronous BTR algorithms. In this section,
 558 we collect the numerical results based on the variant that statically allocates subprob-
 559 lems to the worker processes and chooses the first element of the queue for trial points
 560 (i.e., FIFO). We use the metrics for the performance profile in [6]. We use \mathcal{C} to define
 561 the set of solver configurations. For each problem $p \in \mathcal{P}$ and configuration $c \in \mathcal{C}$, we
 562 use t_{pc} to denote the solution time for solving problem p under configuration c . We
 563 define the performance ratio

564 (48)
$$r_{pc} := \frac{t_{pc}}{\min_{c \in \mathcal{C}} t_{pc}}$$

 565

566 that represents the performance of configuration c as compared with the best per-
 567 formance by any solver configuration on problem p . We now define the performance
 568 of configuration c on any given problem as the probability for configuration c that a





Fig. 3: Distribution of the average percent imbalance metrics resulting from the synchronous BTR method.



(a) All instances

(b) Highly imbalanced instances ($\nu_p \geq 50\%$)

Fig. 4: Performance profile for $\bar{\Lambda} = 1$ and $\underline{\Pi} \in \{1, 4, 8\}$.

569 performance ratio r_{pc} is within a factor τ of the best possible ratio. In other words,

$$570 \quad (49) \quad \rho_c(\tau) := \frac{1}{|\mathcal{P}|} |\{p \in \mathcal{P} : r_{pc} \leq \tau\}|.$$

572 We compare the performance of the synchronous and asynchronous BTR strategies
 573 with different algorithmic settings. In our numerical experiments, we vary the maxi-
 574 mum queue size $\bar{\Lambda} \in \{1, 2\}$ and the minimum number of worker processes to receive
 575 bundle information $\underline{\Pi} \in \{1, 4, 8\}$.

576 **Figure 4** shows the performance of the synchronous and asynchronous BTR algo-
 577 rithms for $\bar{\Lambda} = 1$ and $\underline{\Pi} \in \{1, 4, 8\}$. We label the synchronous method “Sync” and
 578 label the asynchronous method with $\bar{\Lambda} = m$ and $\underline{\Pi} = n$ “Async-Q m P n .” We see
 579 that the asynchronous algorithm results in higher probabilities than the synchronous
 580 counterpart for any factor τ . In **Figure 4a** we present profiles for all instances; we




Fig. 5: Performance profile for $\bar{\Lambda} = 2$ and $\underline{\Pi} \in \{1, 4, 8\}$.

581 can see that Async-Q1P1 has the most wins (with a probability of 0.41) and that
 582 Sync has the least wins (with a probability of 0.11). In Figure 4b we profile the
 583 solvers for highly imbalanced instances ($\mu_p \geq 50\%$); we see that the probability that
 584 Async-Q1P1 is the best solver increases to 0.66, whereas the probability that Sync is
 585 the best solver becomes zero. We also observe that the asynchronous algorithms tend
 586 to be less competitive with a large value of $\underline{\Pi}$. We also note that the asynchronous
 587 algorithm with $\underline{\Pi} = 16$ is equivalent to the synchronous counterpart.

588 In Figure 5 we present results for the case in which we allow for more capacity
 589 for the queue of trial points ($\bar{\Lambda} = 2$). We see that the asynchronous method is faster
 590 than the synchronous method in 87% of the problem instances. Async-Q2P4 is more
 591 competitive than Async-Q2P1 for the highly imbalanced problem instances. Async-
 592 Q2P8 has a lower number of wins than Sync, but the performance becomes much more
 593 competitive if we extend τ of interest to 1.02 or larger. This implies that the more
 594 frequent update of dual variables is not always advantageous from a computational
 595 performance stand-point.

596 **4.5. Variations of Asynchronous Computing.** We present computational
 597 results for the asynchronous BTR algorithm with variations in algorithmic settings.
 598 In particular, we compare different strategies for choosing trial points (i.e., FIFO vs.
 599 LIFO) and for allocating subproblems to worker processes (i.e., static vs. dynamic).
 600 Figures 6 and 7 show the performance plots for the asynchronous algorithms with
 601 the different settings. We label the asynchronous algorithm “X-Y-QmPn” for each
 602 setting $X \in \{Static, Dynamic\}$ and $Y \in \{FIFO, LIFO\}$. To highlight the impact
 603 on performance, we use the highly imbalanced problem instances.

604 Figure 6 shows the performance profiles, as defined in Subsection 4.4, for the
 605 synchronous and asynchronous algorithms with trial points chosen based on FIFO
 606 and LIFO. We perform the numerical experiments with $\bar{\Lambda} = 2$. Note that FIFO and
 607 LIFO are equivalent when $\bar{\Lambda} = 1$. The FIFO policy are faster than the LIFO policy,
 608 regardless of the other algorithmic settings (e.g., static vs. dynamic subproblem
 609 allocations). The reason is that the LIFO policy delays the complete evaluation of
 610 trial points (i.e., satisfying the condition at line 30 in Algorithm 4) by evaluating new
 611 trial points only. In particular, the LIFO policy is slower with the larger queue size




Fig. 6: Performance profile of the asynchronous variations (FIFO vs. LIFO) with $\bar{\Lambda} = 2$ for highly imbalanced instances ($\nu_p \geq 50\%$)

612 (i.e., Q2 vs. Q1).

613 **Figure 7** compares the computational performance for the subproblem allocation
 614 policies (static vs. dynamic). Static allocation is faster than the dynamic allocation,
 615 regardless of the other settings (e.g., FIFO vs. LIFO). The subproblem solution
 616 time must be increased in dynamic allocation, for which the warm-starting feature in
 617 CPLEX is no longer available when different subproblems are solved from iteration to
 618 iteration. We also observe that dynamic allocation is even slower than the synchronous
 619 method for many instances. Consistent to the observations in [Subsection 4.4](#), we found
 620 that the frequent update of dual variables tends to be advantageous, but not always.

621 **4.6. Scalability.** We now demonstrate parallel scalability of the asynchronous
 622 BTR algorithm for the 80 problem instances. We perform the scaling experiments
 623 based on the static subproblem allocation, the FIFO scheme for choosing trial points,
 624 $\bar{\Lambda} = 1$, and $\underline{\Pi} = 1$ (i.e., Static-FIFO-Q1P1). We use 2, 4, 8, 16, and 32 computing
 625 cores to parallelize the asynchronous method, for which half the computing cores
 626 are used in the master problem solution and the others are used in the subproblem
 627 solutions. For the instance with 8 cores, four of the cores are used to parallelize the
 628 subproblem solutions, and the other four are used for solving the master with the
 629 barrier solver. [Figure 8](#) shows scaling performance results of the method. We define
 630 the speedup as the solution time with N cores to that with two cores. The solution
 631 times for each set of 80 instances are shown as a box plot. The linear speedup (red
 632 dashed line in [Figure 8](#)) is achieved when the speedup increases proportional to the
 633 number of cores, which represents the ideal strong scaling efficiency. We observe that
 634 the mean scaling plot for the asynchronous method is closely aligned with the linear
 635 scaling line, which implies that the asynchronous method scales up as the number
 636 of cores increases with respect to the mean solution time. We also highlight that
 637 the scalability results are consistent with those of the synchronous variant reported
 638 in [\[14\]](#). Note, however, that the efficiency degrades as the number of cores increases,
 639 due to Amdahl's law.

640 **5. Summary and Directions of Future Work.** We have developed syn-
 641 chronous and asynchronous variants of a bundle-trust-region (BTR) algorithm within




Fig. 7: Performance profile for the asynchronous variations (Static vs. Dynamic) for highly imbalanced instances ($\nu_p \geq 50\%$)

642 the context of Lagrangian dual decomposition applied to stochastic mixed-integer
 643 programs. The BTR algorithm solves the Lagrangian dual of the SMIP by using a
 644 cutting-plane method with a trust region on the dual search space. In the synchronous
 645 variant, cutting-planes from all scenario subproblems are used to update the trust region
 646 and update the dual search step. We proved that this algorithm converges to the
 647 Lagrangian dual bound of the SMIP and we proved that convergence is independent
 648 of the choice of the trust-region norm and of bundle management steps. Unfortunately,
 649 this algorithm suffers from parallel inefficiencies due to computational load
 650 imbalance in the solution of scenario subproblems (which are solved to obtain the
 651 cutting planes). Motivated by this, we developed an asynchronous variant that uses
 652 only a subset of the subproblem solutions to update the trust region and compute
 653 the dual step. For this method, we devised a trust-region update strategy that uses
 654 only trial points of a queue, while the other trial points may be used to update the
 655 Lagrangian master problem. We also considered the variations of the algorithmic set-
 656 tings: FIFO/LIFO policies for choosing trial points and static/dynamic subproblem
 657 allocations. We proved that all variants of the asynchronous algorithm converge to
 658 the optimal Lagrangian dual bound of the SMIP.




Fig. 8: Scaling efficiency results for asynchronous BTR algorithm (Static-FIFO-Q1P1).

659 The synchronous and asynchronous BTR algorithms are implemented in the open-
 660 source parallel software package DSP. In our numerical experiments, we used the
 661 WECC test system data and created 80 instances of a stochastic unit commitment
 662 problem that schedules a set of power generators and dispatches power to satisfy the
 663 demand of the system under uncertain wind power generation. The results show that
 664 the asynchronous algorithm solves the problem instances significantly faster than the
 665 synchronous counterpart (particularly in the highly imbalanced problem instances).
 666 Moreover, we showed that the asynchronous algorithm achieves strong scaling.

667 If the master problem is relatively more time-consuming than the evaluation of
 668 the subproblems, the computational benefit of the asynchronous approach will not
 669 be as evident. Such a case can be observed when the subproblems are linear pro-
 670 grams and/or when the first stage has a significantly large number of variables. In
 671 such case, the parallelization of master problem solution would improve the computa-
 672 tional performance, as shown in [18]. However, the parallelization approach in [18]
 673 is based on a parallel Schur complement decomposition that would also suffer from the
 674 large number of first-stage variables. Therefore, new parallelization approaches for
 675 the master problem are an interesting research path for future work. In addition, a
 676 computational comparison with other non-smooth methods (e.g., [3, 4, 24]) can be of
 677 interest. As part of future work, we will also seek to improve the asynchronous method
 678 by adaptively changing the parameters $\bar{\Lambda}$ and $\bar{\Pi}$ in order to maximize computational
 679 performance. In particular, highly imbalanced instances can be detected after a few
 680 synchronous iterations, and this information can be used to tune the parameters of
 681 asynchronous iterations. Moreover, inexact evaluation of the Lagrangian subprob-
 682 lems can further be incorporated in the incremental framework to further alleviate
 683 load imbalances (e.g., [16, 7, 24]). Motivated by the observation that the dynamic
 684 allocation is slower than the static allocation as in Subsection 4.5, one can also design
 685 a partial dynamic allocation such that each process can take and solve only certain
 686 subproblems in the allocation scheme, which would allow to use the warm-starting
 687 feature.

688 **Acknowledgments.** This material is based upon work supported by the U.S.
 689 Department of Energy, Office of Science, under contract DE-AC02-06CH11357. The
 690 work of Cosmin G. Petra was performed under the auspices of the U.S. Department
 691 of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-
 692 07NA27344. Victor M. Zavala acknowledges financial support by the U.S. Department
 693 of Energy under grant DE-SC0014114. We gratefully acknowledge the computing
 694 resources provided on Blues, a high-performance computing cluster operated by the
 695 Laboratory Computing Resource Center at Argonne National Laboratory.

696

REFERENCES

- 697 [1] S. AHMED, *A scenario decomposition algorithm for 0–1 stochastic programs*, Operations Research
 698 Letters, 41 (2013), pp. 565–569.
- 699 [2] I. ARAVENA AND A. PAPAVALIOU, *A distributed asynchronous algorithm for the two-stage*
 700 *stochastic unit commitment problem*, in IEEE Power & Energy Society General Meeting,
 701 IEEE, 2015, pp. 1–5.
- 702 [3] D. P. BERTSEKAS, *Incremental proximal methods for large scale convex optimization*, Mathe-
 703 matical Programming, 129 (2011), pp. 163–195.
- 704 [4] D. P. BERTSEKAS, *Incremental aggregated proximal and augmented Lagrangian algorithms*,
 705 arXiv preprint arXiv:1509.09257, (2015).
- 706 [5] C. C. CARØE AND R. SCHULTZ, *Dual decomposition in stochastic integer programming*, Oper-
 707 ations Research Letters, 24 (1999), pp. 37–45.
- 708 [6] E. D. DOLAN AND J. J. MORÉ, *Benchmarking optimization software with performance profiles*,
 709 Mathematical Programming, 91 (2002), pp. 201–213.
- 710 [7] G. EMIEL AND C. SAGASTIZÁBAL, *Incremental-like bundle methods with application to energy*
 711 *planning*, Computational Optimization and Applications, 46 (2010), pp. 305–332.
- 712 [8] F. FISCHER AND C. HELMBERG, *A parallel bundle framework for asynchronous subspace op-*
 713 *timization of nonsmooth convex functions*, SIAM Journal on Optimization, 24 (2014),
 714 pp. 795–822.
- 715 [9] A. FRANGIONI, *Generalized bundle methods*, SIAM Journal on Optimization, 13 (2002), pp. 117–
 716 156.
- 717 [10] M. GAUDIOSO, G. GIALLOMBARDO, AND G. MIGLIONICO, *An incremental method for solving*
 718 *convex finite min-max problems*, Mathematics of Operations Research, 31 (2006), pp. 173–
 719 187.
- 720 [11] K. KIM, *An optimization approach for identifying and prioritizing critical components in a*
 721 *power system*, Tech. Report ANL/MCS-P7076-0717, Argonne National Laboratory, 2017.
- 722 [12] K. KIM, A. BOTTERUD, AND F. QIU, *Temporal decomposition for improved unit commitment in*
 723 *power system production cost modeling*, IEEE Transactions on Power Systems, PP (2018),
 724 pp. 1–1, <https://doi.org/10.1109/TPWRS.2018.2816463>.
- 725 [13] K. KIM, F. YANG, V. M. ZAVALA, AND A. A. CHIEN, *Data centers as dispatchable loads to*
 726 *harness stranded power*, IEEE Transactions on Sustainable Energy, 8 (2017), pp. 208–218.
- 727 [14] K. KIM AND V. M. ZAVALA, *Algorithmic innovations and software for the dual decomposition*
 728 *method applied to stochastic mixed-integer programs*, Mathematical Programming Compu-
 729 tation, (2017), <https://doi.org/10.1007/s12532-017-0128-z>.
- 730 [15] K. C. KIWIEL, *Convergence of approximate and incremental subgradient methods for convex*
 731 *optimization*, SIAM Journal on Optimization, 14 (2004), pp. 807–840.
- 732 [16] K. C. KIWIEL, *A proximal bundle method with approximate subgradient linearizations*, SIAM
 733 Journal on Optimization, 16 (2006), pp. 1007–1023.
- 734 [17] J. LINDEROTH AND S. WRIGHT, *Decomposition algorithms for stochastic programming on a*
 735 *computational grid*, Computational Optimization and Applications, 24 (2003), pp. 207–
 736 250.
- 737 [18] M. LUBIN, K. MARTIN, C. G. PETRA, AND B. SANDIKÇI, *On parallelizing dual decomposition*
 738 *in stochastic integer programming*, Operations Research Letters, 41 (2013), pp. 252–258.
- 739 [19] A. NEDIC AND D. P. BERTSEKAS, *Incremental subgradient methods for nondifferentiable opti-*
 740 *mization*, SIAM Journal on Optimization, 12 (2001), pp. 109–138.
- 741 [20] A. NEDIĆ, D. P. BERTSEKAS, AND V. S. BORKAR, *Distributed asynchronous incremental sub-*
 742 *gradient methods*, Studies in Computational Mathematics, 8 (2001), pp. 381–407.
- 743 [21] A. PAPAVALIOU AND S. S. OREN, *Multiarea stochastic unit commitment for high wind penetra-*
 744 *tion in a transmission constrained network*, Operations Research, 61 (2013), pp. 578–592.
- 745 [22] O. PEARCE, T. GAMBLIN, B. R. DE SUPINSKI, M. SCHULZ, AND N. M. AMATO, *Quantifying*

- 746 *the effectiveness of load balance algorithms*, in Proceedings of the 26th ACM international
747 conference on Supercomputing, ACM, 2012, pp. 185–194.
- 748 [23] K. RYAN, D. RAJAN, AND S. AHMED, *Scenario Decomposition for 0-1 Stochastic Programs:*
749 *Improvements and Asynchronous Implementation*, Optimization Online, (2015).
- 750 [24] W. VAN ACKOOIJ AND A. FRANGIONI, *Incremental bundle methods using upper models*, SIAM
751 Journal on Optimization, 28 (2018), pp. 379–410, <https://doi.org/10.1137/16M1089897>.