

1 Clustering-Based Preconditioning for Stochastic Programs

2 Yankai Cao · Carl D. Laird · Victor M. Zavala

3

4 **Abstract** We present a clustering-based preconditioning strategy for KKT systems arising in
5 stochastic programming within an interior-point framework. The key idea is to perform adap-
6 tive clustering of scenarios (inside-the-solver) based on their influence on the problem at hand.
7 This approach thus contrasts with existing (outside-the-solver) approaches that cluster scenar-
8 ios based on problem data alone. We derive spectral and error properties for the preconditioner
9 and demonstrate that scenario compression rates of up to 87% can be obtained, leading to dra-
10 matic computational savings. In addition, we demonstrate that the proposed preconditioner
11 can avoid scalability issues of Schur decomposition in problems with large first-stage dimen-
12 sionality.

13 **Keywords** preconditioning · interior-point · stochastic · large-scale · clustering.

Preprint Number ANL/MCS-P3050-1112

Yankai Cao · Carl D. Laird
School of Chemical Engineering, Purdue University
480 Stadium Mall Drive, West Lafayette, IN 47907
Tel.: +1-765-494-0085
Fax: +1-765-494-0805
E-mail: {cao142,lairdc}@purdue.edu

Victor M. Zavala
Department of Chemical and Biological Engineering, University of Wisconsin-Madison
1415 Engineering Hall, Madison, WI 53706-1691
Tel.: +1-608-262-6934
Fax: +1-608-262-5434
E-mail: victor.zavala@wisc.edu

14 1 Preliminaries

We consider two-stage stochastic programs of the form

$$\min \left(\frac{1}{2} x_0^T Q_0 x_0 + d_0^T x_0 \right) + \sum_{s \in \mathcal{S}} p_s \left(\frac{1}{2} x_s^T Q_s x_s + d_s^T x_s \right) \quad (1a)$$

$$\text{s.t.} \quad W_0 x_0 = b_0, \quad (y_0) \quad (1b)$$

$$T_s x_0 + W_s x_s = b_s, \quad (y_s), \quad s \in \mathcal{S} \quad (1c)$$

$$x_0 \geq 0, \quad (\nu_0) \quad (1d)$$

$$x_s \geq 0, \quad (\nu_s), \quad s \in \mathcal{S}. \quad (1e)$$

15 Here, $\mathcal{S} := \{1..S\}$ is the scenario set and S is the number of scenarios. The problem variables
 16 are $x_0, \nu_0 \in \mathbb{R}^{n_0}$, $x_s, \nu_s \in \mathbb{R}^{n_s}$, $y_0 \in \mathbb{R}^{m_0}$, and $y_s \in \mathbb{R}^{m_s}$. The total number of variables is
 17 $n := n_0 + \sum_{s \in \mathcal{S}} n_s$, of equality constraints is $m := m_0 + \sum_{s \in \mathcal{S}} m_s$, and of inequalities is n .
 18 We refer to (x_0, y_0, ν_0) as the first-stage variables and to (x_s, y_s, ν_s) , $s \in \mathcal{S}$ as the second-stage
 19 variables. We refer to equation (1a) as the cost function. The *data* defining problem (1) is given by
 20 the cost coefficients d_0, Q_0, Q_s, d_s ; the right-hand side coefficients b_0, b_s ; the matrix coefficients
 21 T_s, W_s ; and the scenario probabilities as $p_s \in \mathbb{R}_+$. We refer to $p_s, Q_s, d_s, b_s, T_s, W_s$ as the *scenario*
 22 *data*.

As is typical in stochastic programming, the number of scenarios can be large and limits the scope of existing off-the-shelf solvers. In this work, we present strategies that cluster scenarios at the linear algebra level to mitigate complexity. We begin by presenting some basic notation. We scale the cost coefficient matrices and vectors as $Q_s \leftarrow p_s Q_s$ and $d_s \leftarrow p_s d_s$. Consequently, the Lagrange function of (1) can be expressed as:

$$\begin{aligned} \mathcal{L}(x, y, \nu) &= \frac{1}{2} x_0^T Q_0 x_0 + d_0^T x_0 + y_0^T (W_0 x_0 - b_0) - \nu_0^T x_0 \\ &\quad + \sum_{s \in \mathcal{S}} \left(\frac{1}{2} x_s^T Q_s x_s + d_s^T x_s + y_s^T (T_s x_0 + W_s x_s - b_s) - \nu_s^T x_s \right). \end{aligned} \quad (2)$$

Here, $x := [x_0^T, x_1^T, \dots, x_S^T]$, $y^T := [y_0^T, y_1^T, \dots, y_S^T]$, and $\nu^T := [\nu_0^T, \nu_1^T, \dots, \nu_S^T]$. In a primal-dual interior-point (IP) setting we seek to solve nonlinear systems of equations of the form

$$\nabla_{x_0} \mathcal{L} = 0 = Q_0 x_0 + d_0 + W_0^T y_0 - \nu_0 + \sum_{s \in \mathcal{S}} T_s^T y_s \quad (3a)$$

$$\nabla_{x_s} \mathcal{L} = 0 = Q_s x_s + d_s + W_s^T y_s - \nu_s, \quad s \in \mathcal{S} \quad (3b)$$

$$\nabla_{y_0} \mathcal{L} = 0 = W_0 x_0 - b_0 \quad (3c)$$

$$\nabla_{y_s} \mathcal{L} = 0 = T_s x_0 + W_s x_s - b_s, \quad s \in \mathcal{S} \quad (3d)$$

$$0 = X_0 V_0 e_{n_0} - \mu e_{n_0} \quad (3e)$$

$$0 = X_s V_s e_{n_s} - \mu e_{n_s}, \quad s \in \mathcal{S}, \quad (3f)$$

with the implicit condition $x_0, \nu_0, x_s, \nu_s \geq 0$. Here, $\mu \geq 0$ is the barrier parameter and $e_{n_0} \in \mathbb{R}^{n_0}$, $e_{n_s} \in \mathbb{R}^{n_s}$ are vectors of ones. We define the diagonal matrices $X_0 := \text{diag}(x_0)$, $X_s := \text{diag}(x_s)$, $V_0 := \text{diag}(\nu_0)$, and $V_s := \text{diag}(\nu_s)$. We define $\alpha_0 := X_0 V_0 e - \mu e_{n_0}$ and $\alpha_s := X_s V_s e -$

μe_{n_s} , $s \in \mathcal{S}$. The search step is obtained by solving the linear system

$$Q_0 \Delta x_0 + W_0^T \Delta y_0 + \sum_{s \in \mathcal{S}} T_s^T \Delta y_s - \Delta \nu_0 = -\nabla_{x_0} \mathcal{L} \quad (4a)$$

$$Q_s \Delta x_s + W_s^T \Delta y_s - \Delta \nu_s = -\nabla_{x_s} \mathcal{L}, \quad s \in \mathcal{S} \quad (4b)$$

$$W_0 \Delta x_0 = -\nabla_{y_0} \mathcal{L} \quad (4c)$$

$$T_s \Delta x_0 + W_s \Delta x_s = -\nabla_{y_s} \mathcal{L}, \quad s \in \mathcal{S} \quad (4d)$$

$$X_0 \Delta \nu_0 + V_0 \Delta x_0 = -\alpha_0 \quad (4e)$$

$$X_s \Delta \nu_s + V_s \Delta x_s = -\alpha_s, \quad s \in \mathcal{S}. \quad (4f)$$

After eliminating the bound multipliers from the linear system we obtain

$$H_0 \Delta x_0 + W_0^T \Delta y_0 + \sum_{s \in \mathcal{S}} T_s^T \Delta y_s = r_{x_0} \quad (5a)$$

$$H_s \Delta x_s + W_s^T \Delta y_s = r_{x_s}, \quad s \in \mathcal{S} \quad (5b)$$

$$W_0 \Delta x_0 = r_{y_0} \quad (5c)$$

$$T_s \Delta x_0 + W_s \Delta x_s = r_{y_s}, \quad s \in \mathcal{S}, \quad (5d)$$

where,

$$H_0 := Q_0 + X_0^{-1} V_0 \quad (6a)$$

$$H_s := Q_s + X_s^{-1} V_s, \quad s \in \mathcal{S}. \quad (6b)$$

We also have that $r_{x_0} := -(\nabla_{x_0} \mathcal{L} - X_0^{-1} \alpha_0)$, $r_{x_s} := -(\nabla_{x_s} \mathcal{L}_s - X_s^{-1} \alpha_s)$, $r_{y_0} := -\nabla_{y_0} \mathcal{L}$, and $r_{y_s} := -\nabla_{y_s} \mathcal{L}$. The step for the bound multipliers can be recovered from

$$\Delta \nu_0 = -X_0^{-1} V_0 \Delta x_0 - X_0^{-1} \alpha_0 \quad (7a)$$

$$\Delta \nu_s = -X_s^{-1} V_s \Delta x_s - X_s^{-1} \alpha_s, \quad s \in \mathcal{S}. \quad (7b)$$

System (5) has the arrowhead form

$$\begin{bmatrix} K_1 & & & B_1 \\ & K_2 & & B_2 \\ & & \ddots & \vdots \\ & & & K_S & B_S \\ B_1^T & B_2^T & \dots & B_S^T & K_0 \end{bmatrix} \begin{bmatrix} \Delta w_1 \\ \Delta w_2 \\ \vdots \\ \Delta w_S \\ \Delta w_0 \end{bmatrix} = \begin{bmatrix} r_1 \\ r_2 \\ \vdots \\ r_S \\ r_0 \end{bmatrix}, \quad (8)$$

where $\Delta w_0^T := [\Delta x_0^T, \Delta y_0^T]$, $\Delta w_s^T := [\Delta x_s^T, \Delta y_s^T]$, $r_0^T := [-r_{x_0}^T, -r_{y_0}^T]$, $r_s^T := [-r_{x_s}^T, -r_{y_s}^T]$, and

$$K_0 := \begin{bmatrix} H_0 & W_0^T \\ W_0 & 0 \end{bmatrix}, \quad K_s := \begin{bmatrix} H_s & W_s^T \\ W_s & 0 \end{bmatrix}, \quad B_s := \begin{bmatrix} 0 & 0 \\ T_s & 0 \end{bmatrix}. \quad (9)$$

²³ We refer to the linear system (8) as the *KKT system* and to its coefficient matrix as the *KKT matrix*. We assume that each scenario block matrix K_s , $s \in \mathcal{S}$ is nonsingular.

²⁴ We use the following notation to define a block-diagonal matrix M composed of blocks M_1, M_2, M_3, \dots :

$$M = \text{blkdiag}\{M_1, M_2, M_3, \dots\}. \quad (10)$$

In addition, we use the following notation to define a matrix B that stacks (row-wise) the blocks B_1, B_2, B_3, \dots :

$$B = \text{rowstack}\{B_1, B_2, B_3, \dots\}. \quad (11)$$

We apply the same rowstack notation for vectors. We use the notation $v^{(k)}$ to indicate the k -th entry of vector v . We use $\text{vec}(M)$ to denote the row-column vectorization of matrix M and we define $\sigma_{\min}(M)$ as the smallest singular value of matrix M . We use $\|\cdot\|$ to denote the Euclidean norm for vectors and the Frobenius norm for matrices, and we recall that $\|M\| = \|\text{vec}(M)\|$ for matrix M .

2 Clustering Setting

In this section we review work on scenario reduction and highlight differences and contributions of our work. We then present our clustering-based preconditioner for the KKT system (8).

2.1 Related Work and Contributions

Scenario clustering (also referred to as aggregation) is a strategy commonly used in stochastic programming to reduce computational complexity. We can classify these strategies as outside-the-solver and inside-the-solver strategies. Outside-the-solver strategies perform clustering on the scenario data (right-hand sides, matrices, and gradients) prior to the solution of the problem [8,16,13,5]. This approach can provide lower bounds and error bounds for linear programs (LPs) and this feature can be exploited in branch-and-bound procedures [5,1,22,28].

Outside-the-solver clustering approaches give rise to several inefficiencies. First, several optimization problems might need to be solved in order to refine the solution. Second, these approaches focus on the problem data and thus *do not capture the effect of the data on the particular problem at hand*. Consider, for instance, the situation in which the same scenario data (e.g., weather scenarios) is used for two different problem classes (e.g., farm planning and power grid planning). Moreover, clustering scenarios based on data alone is inefficient because scenarios that are close in terms of data might have very different impact on the cost function (e.g., if they are close to the constraint boundary). Conversely, two scenarios that are distant in terms of data might have similar contributions to the cost function. We also highlight that many scenario generation procedures require knowledge of the underlying probability distributions [10,13] which are often not available in closed form (e.g., weather forecasting) [26,18].

In this work, we seek to overcome these inefficiencies by performing clustering adaptively inside-the-solver. In an interior-point setting this can be done by creating a preconditioner for the KKT system (8) by *clustering* the scenario blocks. A key advantage of this approach is that a single optimization problem is solved and the clusters are refined only if the preconditioner is not sufficiently accurate. In addition, this approach provides a mechanism to capture the influence of the data on the particular problem at hand. Another advantage is that it can enable *sparse preconditioning* of Schur complement systems. This is beneficial in situations where the number of first-stage variables is large and thus Schur complement decomposition is expensive. Moreover, our approach does not require any knowledge of the underlying probability distributions generating the scenario data. Thus, it can be applied to problems in which simulators are used to generate scenarios (e.g., weather forecasting), and it can be applied to problem classes

63 that exhibit similar structures such as support vector machines [11,14] and scenario-based ro-
 64 bust optimization [4]. Our proposed clustering approach can also be used in combination with
 65 outside-the-solver scenario aggregation procedures, if desired.

66 Related work on inside-the-solver scenario reduction strategies includes stochastic Newton
 67 methods [3]. These approaches sample scenarios to create a smaller representation of the KKT
 68 system. Existing approaches, however, cannot handle constraints. Scenario and constraint re-
 69 duction approaches for IP solvers have been presented in [6,24,20,7]. In [24], scenarios that
 70 have little influence on the step computation are eliminated from the optimality system. This
 71 influence is measured in terms of the magnitude of the constraint multipliers or in terms of the
 72 products $X_s^{-1}V_s$. In that work, it was found that a large proportion of scenarios or constraints
 73 can be eliminated without compromising convergence. The elimination potential can be lim-
 74 ited in early iterations, however, because it is not clear which scenarios have strong or weak
 75 influence on the solution. In addition, this approach eliminates the scenarios from the problem
 76 formulation, and thus special safeguards are needed to guarantee convergence. Our proposed
 77 clustering approach does not eliminate the scenarios from the problem formulation; instead, the
 78 scenario space is compressed to construct preconditioners.

79 In [20] preconditioners for Schur systems are constructed by sampling the full scenario set.
 80 A shortcoming of this approach is that scenario outliers with strong influence might not be cap-
 81 tured in the preconditioner. This behavior is handled more efficiently in the preconditioner pro-
 82 posed in [6] in which scenarios having strong influence on the Schur complement are retained
 83 and those that have weak influence are eliminated. A limitation of the Schur preconditioners
 84 proposed in [20,6] is that they require a dense preconditioner for the Schur complement, which
 85 hinders scalability in problems with many first-stage variables. Our preconditioning approach
 86 enables sparse preconditioning and thus avoids forming and factorizing dense Schur comple-
 87 ments. In addition, compared with approaches in [6,24,20], our approach clusters scenarios
 88 instead of eliminating them (either by sampling or by measuring strong/weak influence). This
 89 enables us to handle scenario redundancies and outliers. In [7], scenarios are clustered to solve
 90 a reduced problem and the solution of this problem is used to warm-start the problem defined
 91 for the full scenario set. The approach can reduce the number of iterations of the full scenario
 92 problem; but the work per iteration is not reduced, as in our approach.

93 2.2 Clustering-Based Preconditioner

To derive our clustering-based preconditioner, we partition the full scenario set \mathcal{S} into C clus-
 ters, where $C \leq S$. We define the cluster set $\mathcal{C} := \{1, \dots, C\}$ and a partition of the scenario set
 $\{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_C\}$ with $\omega_i := |\mathcal{S}_i|$, $i \in \mathcal{C}$. That is,

$$\bigcup_{i \in \mathcal{C}} \mathcal{S}_i = \mathcal{S} \quad (12a)$$

$$\mathcal{S}_i \cap \mathcal{S}_j = \emptyset, \quad i, j \in \mathcal{C}, j \neq i. \quad (12b)$$

For each cluster $i \in \mathcal{C}$, we pick an index $c_i \in \mathcal{S}_i$ to represent the cluster and we use these indexes
 to define the compressed set $\mathcal{R} := \{c_1, c_2, \dots, c_C\}$. We note that $|\mathcal{R}| = C$ and that ω_i are cluster
 weights. We define the binary indicator $\kappa_{s,i}$, $s \in \mathcal{S}$, $i \in \mathcal{C}$, satisfying

$$\kappa_{s,i} = \begin{cases} 1 & \text{if } s \in \mathcal{S}_i \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

Using this notation we have that for arbitrary vectors $v_{c_i}, v_s, i \in \mathcal{C}$, the following identities hold:

$$\sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} \|v_{c_i} - v_s\| = \sum_{s \in \mathcal{S}} \sum_{i \in \mathcal{C}} \kappa_{s,i} \|v_{c_i} - v_s\| \quad (14a)$$

$$\sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} v_s = \sum_{s \in \mathcal{S}} v_s \quad (14b)$$

$$\sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} v_{c_i} = \sum_{i \in \mathcal{C}} \omega_i v_{c_i}. \quad (14c)$$

94 At this point, we have yet to define appropriate procedures for obtaining the cluster information
95 $\mathcal{R}, \mathcal{S}_i, \omega_i$ and $\kappa_{s,i}$. These will be discussed in Section 3.

Consider now the compact representation of the KKT system (8),

$$\underbrace{\begin{bmatrix} K_S & B_S \\ B_S^T & K_0 \end{bmatrix}}_{:=K} \underbrace{\begin{bmatrix} q_S \\ q_0 \end{bmatrix}}_{:=q} = \underbrace{\begin{bmatrix} t_S \\ t_0 \end{bmatrix}}_{:=t}, \quad (15)$$

where

$$K_S := \text{blkdiag} \{K_1, \dots, K_S\} \quad (16a)$$

$$B_S := \text{rowstack} \{B_1, \dots, B_S\} \quad (16b)$$

$$q_S := \text{rowstack} \{q_1, \dots, q_S\} \quad (16c)$$

$$t_S := \text{rowstack} \{t_1, \dots, t_S\}. \quad (16d)$$

Here, (t_0, t_S) are arbitrary right-hand side vectors and (q_0, q_S) are solution vectors.¹ If the solution vector (q_0, q_S) does not exactly solve (15), it will induce a residual vector that we define as $\epsilon_r^T := [\epsilon_{r_0}^T, \epsilon_{r_S}^T]$ with

$$\epsilon_{r_0} := K_0 q_0 + B_S^T q_S - t_0 \quad (17a)$$

$$\epsilon_{r_S} := K_S q_S + B_S q_0 - t_S. \quad (17b)$$

The Schur system of (15) is given by

$$\underbrace{(K_0 - B_S^T K_S^{-1} B_S)}_{:=Z} q_0 = \underbrace{t_0 - B_S^T K_S^{-1} t_S}_{:=t_Z}. \quad (18)$$

Because K_S is block-diagonal, we have that

$$Z = K_0 - \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} B_s^T K_s^{-1} B_s \quad (19a)$$

$$t_Z = t_0 - \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} B_s^T K_s^{-1} t_s. \quad (19b)$$

We now define the following:

$$K_{\mathcal{R}}^\omega := \text{blkdiag} \{\omega_1 K_{c_1}, \omega_2 K_{c_2}, \dots, \omega_C K_{c_C}\} \quad (20a)$$

$$K_{\mathcal{R}}^{1/\omega} := \text{blkdiag} \{1/\omega_1 K_{c_1}, 1/\omega_2 K_{c_2}, \dots, 1/\omega_C K_{c_C}\} \quad (20b)$$

$$B_{\mathcal{R}} := \text{rowstack} \{B_{c_1}, B_{c_2}, \dots, B_{c_C}\} \quad (20c)$$

$$t_{\mathcal{R}} := \text{rowstack} \{t_{c_1}, t_{c_2}, \dots, t_{c_C}\}. \quad (20d)$$

¹ We make a slight remark regarding notation: K_S is a block diagonal matrix while K_S in an entry of such block matrix. A similar observation applies to matrices B_S and vectors q_S, t_S with corresponding entries B_s, q_s, t_s .

In other words, $K_{\mathcal{R}}^{\omega}$ is a block-diagonal matrix in which each block entry K_{c_i} is weighted by the scalar weight ω_i and $K_{\mathcal{R}}^{1/\omega}$ is a block-diagonal matrix in which each block entry K_{c_i} is weighted by $1/\omega_i$. Note that

$$(K_{\mathcal{R}}^{1/\omega})^{-1} = (K_{\mathcal{R}}^{-1})^{\omega}, \quad (21)$$

where,

$$(K_{\mathcal{R}}^{-1})^{\omega} := \text{blkdiag} \{ \omega_1 K_{c_1}^{-1}, \omega_2 K_{c_2}^{-1}, \dots, \omega_C K_{c_C}^{-1} \}. \quad (22)$$

We now present the *clustering-based preconditioner* (CP),

$$\begin{bmatrix} K_{\mathcal{R}}^{1/\omega} & B_{\mathcal{R}} \\ B_{\mathcal{R}}^T & K_0 \end{bmatrix} \begin{bmatrix} \cdot \\ q_0 \end{bmatrix} = \begin{bmatrix} t_{\mathcal{R}} \\ t_0 + t_{CP} \end{bmatrix} \quad (23a)$$

$$K_s q_s = t_s - B_s q_0, \quad i \in \mathcal{C}, s \in \mathcal{S}_i, \quad (23b)$$

where

$$t_{CP} := \sum_{i \in \mathcal{C}} \omega_i B_{c_i}^T K_{c_i}^{-1} t_{c_i} - \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} B_s^T K_s^{-1} t_s \quad (24)$$

is a correction term that is used to achieve consistency between CP and the KKT system. In particular, the Schur system of (23a) is

$$\begin{aligned} \bar{Z} q_0 &= t_0 + t_{CP} - B_{\mathcal{R}}^T (K_{\mathcal{R}}^{1/\omega})^{-1} t_{\mathcal{R}} \\ &= t_0 + t_{CP} - \sum_{i \in \mathcal{C}} \omega_i B_{c_i}^T K_{c_i}^{-1} t_{c_i} \\ &= t_0 - \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} B_s^T K_s^{-1} t_s \\ &= t_Z, \end{aligned} \quad (25)$$

with

$$\begin{aligned} \bar{Z} &:= K_0 - \sum_{i \in \mathcal{C}} \omega_i B_{c_i}^T K_{c_i}^{-1} B_{c_i} \\ &= K_0 - \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} B_{c_i}^T K_{c_i}^{-1} B_{c_i}. \end{aligned} \quad (26)$$

Consequently, the Schur system of the preconditioner and of the KKT system have the same right-hand side. This property is key to establish spectral and error properties for the preconditioner. In particular, note that the solution of CP system (23a)-(23b) solves the perturbed KKT system,

$$\underbrace{\begin{bmatrix} K_{\mathcal{S}} & B_{\mathcal{S}} \\ B_{\mathcal{S}}^T & K_0 + E_Z \end{bmatrix}}_{:= \bar{K}} \begin{bmatrix} q_{\mathcal{S}} \\ q_0 \end{bmatrix} = \begin{bmatrix} t_{\mathcal{S}} \\ t_0 \end{bmatrix}, \quad (27)$$

where

$$E_Z := \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} B_s^T K_s^{-1} B_s - \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} B_{c_i}^T K_{c_i}^{-1} B_{c_i}, \quad (28)$$

is the Schur error matrix and satisfies $Z + E_Z = \bar{Z}$. The mathematical equivalence between CP system (23a)-(23b) and (27) can be established by constructing the Schur system of (27) and noticing that it is equivalent to (25). Moreover, the second-stage steps are the same. Consequently, *applying preconditioner CP is equivalent to using the perturbed matrix \bar{K} as a preconditioning matrix for the KKT matrix K .*

The main idea behind preconditioner CP (we will use CP for short) is to compress the KKT system (15) into the smaller system (23a) which is cheaper to factorize. We solve this smaller system to obtain q_0 , and we recover q_S from (23b) by factorizing the individual blocks K_s . We refer to the coefficient matrix of (23a) as the *compressed matrix*.

In the following, we assume that the Schur complements Z and \bar{Z} are nonsingular. The nonsingularity of Z together with the assumption that all the blocks K_s are nonsingular implies (from the Schur complement theorem) that matrix K defined in (15) is nonsingular and thus the KKT system has a unique solution. The nonsingularity of \bar{Z} together with the assumption that all the blocks K_s are nonsingular implies that the compressed matrix is nonsingular and thus CP has a unique solution. Note that we could have also assumed nonsingularity of matrix K directly and this, together with the nonsingularity of the blocks K_s , would imply nonsingularity of Z (this also from the Schur complement theorem). The same applies if we assume nonsingularity of the compressed matrix, which would imply nonsingularity of \bar{Z} .

Schur decomposition is a popular approach for solving structured KKT systems on parallel computers but it suffers from poor scalability with the dimension of q_0 . The reason is that the Schur complement needs to be formed (this requires as many backsolves with the factors of K_s as the dimension of q_0) and factored (this requires a factorization of a nearly dense matrix of dimension q_0). We elaborate on these scalability issues in Section 4. We thus highlight that the Schur system representations are used only for analyzing CP.

Our preconditioning setting is summarized as follows. At each IP iteration k , we compute a step by solving the KKT system (8). We do so by finding a solution vector $(\Delta w_0, \Delta w_S)$ of the ordered KKT system (8) for the right-hand side (r_0, r_S) using an iterative linear algebra solver such as GMRES, QMR, or BICGSTAB. Here, (r_0, r_S) are the right-hand side vectors of the KKT system (8) in ordered form. Each minor iteration of the iterative linear algebra solver is denoted by $\ell = 0, 1, 2, \dots$. We denote the initial guess of the solution vector of (8) as $(\Delta w_0^\ell, \Delta w_S^\ell)$ with $\ell = 0$. At each minor iterate ℓ , the iterative solver will request the application of CP to a given vector (t_0^ℓ, t_S^ℓ) , and the solution vectors (q_0^ℓ, q_S^ℓ) of (23) are returned to the iterative linear algebra solver. *Perfect preconditioning* occurs when we solve (8) instead of (23) with the right-hand sides (t_0^ℓ, t_S^ℓ) .

3 Preconditioner Properties

In this section we establish properties for CP and we use these to guide the design of appropriate clustering strategies. Because the CP system (23) and the perturbed KKT system (27) are equivalent, we can establish the following result.

Lemma 1 *The preconditioned matrix $\bar{K}^{-1}K$ has $(n+m-n_0-m_0)$ unit eigenvalues, and the remaining (n_0+m_0) eigenvalues are bounded as*

$$|\lambda(\bar{K}^{-1}K) - 1| \leq \frac{1}{\sigma_{\min}(\bar{Z})} \|E_Z\|.$$

Proof: The eigenvalues λ and eigenvectors $w := (w_S, w_0)$ of $\bar{K}^{-1}K$ satisfy $\bar{K}^{-1}Kw = \lambda w$, and thus $Kw = \lambda\bar{K}w$. Consequently,

$$\begin{aligned} K_S w_S + B_S w_0 &= \lambda(K_S w_S + B_S w_0) \\ B_S^T w_S + K_0 w_0 &= \lambda B_S^T w_S + \lambda(K_0 + E_Z)w_0. \end{aligned}$$

From the first relationship we have $n + m - n_0 - m_0$ unit eigenvalues. Applying Schur decomposition to the eigenvalue system, we obtain

$$\begin{aligned} Z w_0 &= \lambda(Z + E_Z)w_0 \\ &= \lambda\bar{Z}w_0. \end{aligned}$$

We can thus express the remaining $n_0 + m_0$ eigenvalues of $\bar{K}^{-1}K$ as $\lambda = 1 + \epsilon_Z$ to obtain

$$\begin{aligned} |\epsilon_Z| &= \frac{\|E_Z w_0\|}{\|\bar{Z}w_0\|} \\ &\leq \frac{1}{\sigma_{\min}(\bar{Z})} \|E_Z\|. \end{aligned}$$

The proof is complete. \square

The above lemma is a direct consequence of Theorem 3.1 in [9]. From the definition of E_Z we note that the following bound holds:

$$|\lambda(\bar{K}^{-1}K) - 1| \leq \frac{1}{\sigma_{\min}(\bar{Z})} \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} \|B_s^T K_s^{-1} B_s - B_{c_i}^T K_{c_i}^{-1} B_{c_i}\|. \quad (30)$$

Lemma 1 states that we can improve the spectrum of $\bar{K}^{-1}K$ by choosing clusters that minimize $\|E_Z\|$. This approach, however, would require expensive matrix operations. An interesting and tractable exception occurs when $Q_s = Q$, $W_s = W$, and $T_s = T$, $i \in \mathcal{C}$, $s \in \mathcal{S}_i$. This case is quite common in applications and arises when the scenario data is only defined by the right-hand sides b_s and the cost coefficients d_s of (1). We refer to this case as the *special data case*. In this case we have that E_Z reduces to

$$E_Z = \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} B^T (K_s^{-1} - K_{c_i}^{-1}) B. \quad (31)$$

We also have that K_s and K_{c_i} differ only in the diagonal matrices $X_s^{-1}V_s$ and $X_{c_i}^{-1}V_{c_i}$. We thus have,

$$K_s - K_{c_i} = \begin{bmatrix} (X_s^{-1}V_s - X_{c_i}^{-1}V_{c_i}) & 0 \\ 0 & 0 \end{bmatrix}. \quad (32)$$

If we define the vectors,

$$\gamma_s = \text{vec}(X_s^{-1}V_s), i \in \mathcal{C}, s \in \mathcal{S}_i \quad (33a)$$

$$\gamma_{c_i} = \text{vec}(X_{c_i}^{-1}V_{c_i}), i \in \mathcal{C}, \quad (33b)$$

¹³⁴ we can establish the following result.

Theorem 1 Assume that $Q_s = Q$, $W_s = W$, and $T_s = T$, $i \in \mathcal{C}$, $s \in \mathcal{S}_i$ holds. Let vectors γ_s, γ_{c_i} be defined as in (33). The preconditioned matrix $\bar{K}^{-1}K$ has $(n + m - n_0 - m_0)$ unit eigenvalues, and there exists a constant $c_K > 0$ such that the remaining $(n_0 + m_0)$ eigenvalues are bounded as

$$|\lambda(\bar{K}^{-1}K) - 1| \leq \frac{c_K}{\sigma_{\min}(\bar{Z})} \sum_{s \in \mathcal{S}} \sum_{i \in \mathcal{C}} \kappa_{s,i} \|\gamma_{c_i} - \gamma_s\|.$$

Proof: From Lemma 1 we have that $n_0 + m_0$ eigenvalues λ of $\bar{K}^{-1}K$ are bounded as $|\lambda - 1| \leq \frac{1}{\sigma_{\min}(Z)} \|E_Z\|$. We define the error matrix,

$$E_s := K_s - K_{c_i}, \quad i \in \mathcal{C}, s \in \mathcal{S}_i$$

and use (31) and (32) to obtain the bound,

$$\begin{aligned} \|E_Z\| &\leq \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} \|B^T B\| \|K_s^{-1} - K_{c_i}^{-1}\| \\ &= \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} \|B^T B\| \|(K_{c_i} + E_s)^{-1} - K_{c_i}^{-1}\|. \end{aligned}$$

We have that

$$\begin{aligned} (K_{c_i} + E_s)^{-1} - K_{c_i}^{-1} &= -(K_{c_i} + E_s)^{-1} E_s K_{c_i}^{-1} \\ &= -K_s^{-1} E_s K_{c_i}^{-1}. \end{aligned}$$

This can be verified by multiplying both sides by $K_{c_i} + E_s$. We thus have

$$\begin{aligned} \|E_Z\| &\leq \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} \|B^T B\| \|(K_{c_i} + E_s)^{-1} - K_{c_i}^{-1}\| \\ &\leq \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} \|B^T B\| \|K_s^{-1}\| \|K_{c_i}^{-1}\| \|E_s\| \\ &\leq c_K \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} \|\text{vec}(X_{c_i}^{-1} V_{c_i}) - \text{vec}(X_s^{-1} V_s)\|, \end{aligned}$$

135 with $c_K := \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} \|B^T B\| \|K_s^{-1}\| \|K_{c_i}^{-1}\|$. The existence of c_K follows from the nonsingularity of K_s and K_{c_i} . The proof is complete. \square

137

138 We now develop a bound of the preconditioning error for the *general data case* in which
139 the scenario data is also defined by coefficient matrices. Notably, this bound does not require
140 the minimization of the error $\|E_Z\|$. The idea is to bound the error induced by CP relative to
141 the exact solution of the KKT system (15) (perfect preconditioner). This approach is used to
142 characterize inexact preconditioners such as multigrid and nested preconditioned conjugate
143 gradient [23]. We express the solution of CP obtained from (23) as $q^T = [q_S^T, q_0^T]$ and that of the
144 KKT system (15) as $q^{*T} = [q_S^{*T}, q_0^{*T}]$. We define the error between q and q^* as $\epsilon := q - q^*$ and
145 we seek to bound ϵ . If we decompose the error as $\epsilon^T = [\epsilon_S^T, \epsilon_0^T]$, we have that $\epsilon_0 = q_0 - q_0^*$ and
146 $\epsilon_S = q_S - q_S^*$.

We recall that the Schur systems of (15) and of (23) and their respective solutions satisfy

$$Zq_0^* = t_Z \tag{34a}$$

$$\bar{Z}q_0 = t_Z. \tag{34b}$$

If we define the vectors,

$$\gamma_s = (B_s^T K_s^{-1} B_s) t_Z, \quad i \in \mathcal{C}, s \in \mathcal{S}_i \tag{35a}$$

$$\gamma_{c_i} = (B_{c_i}^T K_{c_i}^{-1} B_{c_i}) t_Z, \quad i \in \mathcal{C}. \tag{35b}$$

147 we can establish the following bound on the error $\epsilon = q - q^*$.

Lemma 2 Assume that there exists $c_T > 0$ such that $\|(Z - \bar{Z})Z^{-1}t_Z\| \leq c_T\|(Z - \bar{Z})t_Z\|$ holds; then there exists $c_{ZK} > 0$ such that the preconditioner error ϵ is bounded as

$$\|\epsilon\| \leq c_{ZK}\|(Z - \bar{Z})t_Z\|.$$

Proof: From $\epsilon_0 = q_0 - q_0^*$ we have $\bar{Z}\epsilon_0 = \bar{Z}q_0 - \bar{Z}q_0^*$. From (34) we have $\bar{Z}q_0 = Zq_0^* = t_Z$ and thus $\bar{Z}\epsilon_0 = Zq_0^* - \bar{Z}q_0^*$. We thus have,

$$\begin{aligned} \bar{Z}\epsilon_0 &= Zq_0^* - \bar{Z}q_0^* \\ &= t_Z - \bar{Z}q_0^* \\ &= t_Z - \bar{Z}Z^{-1}t_Z \\ &= t_Z - (Z + (\bar{Z} - Z))Z^{-1}t_Z \\ &= (Z - \bar{Z})Z^{-1}t_Z. \end{aligned}$$

We recall that

$$\begin{aligned} q_S^* &= K_S^{-1}(t_S - B_S q_0^*) \\ q_S &= K_S^{-1}(t_S - B_S q_0) \end{aligned}$$

and thus

$$\begin{aligned} \epsilon_S &= K_S^{-1}B_S(q_0^* - q_0) \\ &= -K_S^{-1}B_S\epsilon_0. \end{aligned}$$

We thus have

$$\begin{aligned} \|\epsilon_0\| &\leq c_Z\|(Z - \bar{Z})t_Z\| \\ \|\epsilon_S\| &\leq c_{K_S}\|\epsilon_0\|, \end{aligned}$$

148 with $c_Z := \|\bar{Z}^{-1}\|c_T$ and $c_{K_S} := \|K_S^{-1}B_S\|$. The existence of c_Z follows from the assumption
149 that \bar{Z} is nonsingular. The existence of c_{K_S} follows from the assumption that the blocks K_s are
150 nonsingular and thus K_S is nonsingular. The result follows from $\|\epsilon\| \leq \|\epsilon_0\| + \|\epsilon_S\|$ and by
151 defining $c_{ZK} := c_Z(1 + c_{K_S})$. \square

152

153 The assumption that there exists $c_T > 0$ such that $\|(Z - \bar{Z})Z^{-1}t_Z\| \leq c_T\|(Z - \bar{Z})t_Z\|$ holds
154 is trivially satisfied when Z^{-1} and \bar{Z} commute (i.e., $\bar{Z}Z^{-1}$ is a symmetric matrix). In this case
155 we have that $c_T = \|Z^{-1}\|$. The matrices also commute in the limit $\bar{Z} \rightarrow Z$ because $\bar{Z}Z^{-1} =$
156 $ZZ^{-1} + (\bar{Z} - Z)Z^{-1}$ and thus $\bar{Z}Z^{-1} \rightarrow I$. When Z and \bar{Z} do not commute we require that
157 $\|(Z - \bar{Z})Z^{-1}t_Z\|$ decreases when $\|(Z - \bar{Z})t_Z\|$ does. We validate this condition empirically in
158 Section 4.

Theorem 2 Let vectors γ_s, γ_{c_i} be defined as in (35). The preconditioner error ϵ is bounded as

$$\|\epsilon\| \leq c_{ZK} \sum_{s \in \mathcal{S}} \sum_{i \in \mathcal{C}} k_{s,i} \|\gamma_{c_i} - \gamma_s\|,$$

159 with c_{ZK} defined in Lemma 2.

Proof: From (35) and (28) we have that

$$\begin{aligned}
\bar{Z}t_Z - Zt_Z &= E_Z t_Z \\
&= \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} B_s^T K_s^{-1} B_s t_Z - \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} B_{c_i}^T K_{c_i}^{-1} B_{c_i} t_Z \\
&= \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} (B_s^T K_s^{-1} B_s t_Z - B_{c_i}^T K_{c_i}^{-1} B_{c_i} t_Z) \\
&= \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} (\gamma_s - \gamma_{c_i}).
\end{aligned}$$

We bound this expression to obtain,

$$\begin{aligned}
\|\bar{Z}t_Z - Zt_Z\| &= \left\| \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} (\gamma_s - \gamma_{c_i}) \right\| \\
&\leq \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} \|\gamma_{c_i} - \gamma_s\| \\
&= \sum_{s \in \mathcal{S}} \sum_{i \in \mathcal{C}} \kappa_{s,i} \|\gamma_{c_i} - \gamma_s\|.
\end{aligned}$$

160 The result follows from Lemma 2. \square

161

We can see that the properties of CP are related to a metric of the form

$$\mathcal{D}_C := \sum_{s \in \mathcal{S}} \sum_{i \in \mathcal{C}} \kappa_{s,i} \|\gamma_{c_i} - \gamma_s\|. \quad (36)$$

162 This is the *distortion metric* widely used in clustering analysis [2]. The distortion metric is (par-
163 tially) minimized by K-means, K-medoids, and hierarchical clustering algorithms to determine
164 $\kappa_{s,i}$ and γ_{c_i} . The vectors γ_s are called *features*, and γ_{c_i} is the centroid of cluster $i \in \mathcal{C}$ (we can also
165 pick the scenario that is closest to the centroid if the centroid is not an element of the scenario
166 set). The distortion metric is interpreted as the accumulated distance of the elements of the cluster
167 relative to the centroid. If the distortion is small, then the scenarios in a cluster are similar.
168 The distortion metric can be made arbitrarily small by increasing the number of clusters and is
169 zero in the limit with $S = C$ because each cluster is given by one scenario. Consequently, we see
170 that Theorems 1 and 2 provide the necessary insights to derive clusters using different sources
171 of information of the scenarios.

172 Theorem 1 suggests that, in the special data case with features defined as $\gamma_s = \text{vec}(X_s^{-1}V_s)$,
173 the spectrum of $K^{-1}K$ can be made arbitrarily close to one if the distortion metric is made ar-
174 bitrarily small. This implies that the *definition of the features is consistent*. We highlight, however,
175 that the bounds of Theorem 1 assume that the clustering parameters are given (i.e., the sets \mathcal{C}
176 and \mathcal{C}_i are fixed). Consequently, the constants c_K , and $\sigma_{\min}(\bar{Z})$ change when the clusters are
177 changed. Because of this, we cannot guarantee that reducing the distortion metric will indeed
178 improve the quality of the preconditioner. The aforementioned constants depend in nontrivial
179 ways on the clustering parameters and it is thus difficult to obtain bounds for them. In the next
180 section we demonstrate empirically, however, that the constants c_K and $\sigma_{\min}(\bar{Z})$ are insensitive
181 to the clustering parameters. Consequently, reducing the distortion metric in fact improves the
182 quality of the preconditioner. We leave the theoretical treatment of this issue as part of future
183 work.

184 We can obtain useful insights from the special data case. First note that the scenarios are
 185 clustered at each IP iteration k because the matrices $X_s^{-1}V_s$ change along the search. The clus-
 186 tering approach is therefore adaptive, unlike outside-the-solver scenario clustering approaches.
 187 In fact, it is not possible to derive spectral and error properties for preconditioners based on
 188 clustering of problem data alone. Our approach focuses directly on the contributions $X_s^{-1}V_s$
 189 and thus assumes that the problem data enters indirectly through the contributions $X_s^{-1}V_s$,
 190 which in turn affect the structural properties of the KKT matrix. The features $\gamma_s = \text{vec}(X_s^{-1}V_s)$
 191 have an important interpretation: these reflect the contribution of each scenario to the logarithmic
 192 barrier function. From complementarity we have that $\|X_s\| \gg 0$ implies $\|V_s\| \approx 0$ and
 193 $\|X_s^{-1}V_s\| \approx 0$. In this case we say that there is weak activity in the scenario and we have from
 194 (6) that $H_s = Q_s + X_s^{-1}V_s \approx Q_s$. Consequently, the primal-dual term $X_s^{-1}V_s$ for a scenario with
 195 weak activity puts little weight on the barrier function. In the opposite case in which the sce-
 196 nario has strong activity we have that $\|V_s\| \gg 0$, $\|X_s\| \approx 0$, and $\|X_s^{-1}V_s\| \gg 0$. In this case we
 197 thus have that a scenario with strong activity puts a large weight on the barrier function. This
 198 reasoning is used in [24, 12] to eliminate the scenarios with weak activity. In our case we pro-
 199 pose to cluster scenarios with similar activities. Clustering allows us to eliminate redundancies
 200 in both active and inactive scenarios and to capture outliers. In addition, this strategy avoids
 201 the need to specify a threshold to classify weak and strong activity.

202 Theorem 2 provides a mechanism to obtain clusters for the general data case in which the
 203 scenario data is defined also by the coefficient matrices. The result states that we can bound
 204 the preconditioning error using the Schur complement error $E_Z = \bar{Z} - Z$ projected on the
 205 right-hand side vector t_Z . Consequently, the error can be bounded by the distortion metric with
 206 features defined in (35). This suggests that the error can be made arbitrarily small if the distor-
 207 tion is made arbitrarily small. Moreover, it is not necessary to perform major matrix operations.
 208 As in the special data case of Theorem 1, however, the bounding constant c_Z of Theorem 2 de-
 209 pends on the clustering parameters. Moreover, we need to verify that the term $\|(Z - \bar{Z})Z^{-1}t_Z\|$
 210 decreases when $\|(Z - \bar{Z})t_Z\|$ does. In the next section we verify these two assumptions empiri-
 211 cally.

The error bound of Theorem 2 requires that clustering tasks and the factorization of the com-
 pressed matrix be performed at each minor iteration ℓ of the iterative linear algebra solver. The
 reason is that the features (35) change with t_Z^ℓ . Performing these tasks at each minor iteration,
 however, is expensive. Consequently, we perform these tasks only at the first minor iteration
 $\ell = 0$. If the initial guess of the solution vector of the KKT system is set to zero ($\Delta w_0^\ell = 0$ and
 $\Delta w_S^\ell = 0$) and if GMRES, QMR, or BICGSTAB schemes are used, this is equivalent to perform-
 ing by clustering using the features

$$\gamma_s = (B_s^T K_s^{-1} B_s) r_Z, i \in \mathcal{C}, s \in \mathcal{S}_i \quad (37a)$$

$$\gamma_{c_i} = (B_{c_i}^T K_{c_i}^{-1} B_{c_i}) r_Z, i \in \mathcal{C}, \quad (37b)$$

where

$$\begin{aligned} r_Z &= t_Z^0 \\ &= r_0 - \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} B_s^T K_s^{-1} r_s \end{aligned} \quad (38)$$

212 is the right-hand side of the Schur system of (8).

213 4 Numerical Results

214 In this section we discuss implementation issues of CP and present numerical results for bench-
 215 mark problems in the literature and a large-scale stochastic market clearing problem. We begin
 216 by summarizing the procedure for computing the step $(\Delta x_k, \Delta y_k, \Delta \nu_k)$ at each IP iteration k .

217 Step Computation Scheme

- 219 1. **Initialization.** Given iterate (x_k, y_k, ν_k) , number of clusters C , tolerance τ_k , and maximum
 220 number of linear solver iterates m_{it} .
- 221 2. **Get Clustering Information.**
- 222 2.0. Compute features γ_s , $s \in \mathcal{S}$ as in (33) or (37).
- 223 2.1. Obtain $\kappa_{s,i}$ and γ_{c_i} using a clustering algorithm (e.g., K-means, hierarchical).
- 224 2.2. Use $\kappa_{s,i}$ to construct \mathcal{C} , \mathcal{R} , and ω_i .
- 225 2.3. Construct and factorize compressed matrix

$$\begin{bmatrix} K_{\mathcal{R}}^{1/\omega} & B_{\mathcal{R}} \\ B_{\mathcal{R}}^T & K_0 \end{bmatrix}$$

225 and factorize scenario matrices K_s , $i \in \mathcal{C}$, $s \in \mathcal{S}_i$.

226 3. Get Step.

- 227 3.1. Call iterative linear solver to solve KKT system (15) with right-hand sides $(r_0, r_{\mathcal{S}})$, set
 228 $\ell = 0$, and initial guess $\Delta w_0^\ell = 0$ and $\Delta w_{\mathcal{S}}^\ell = 0$. At each minor iterate $\ell = 0, 1, \dots$, of the
 229 iterative linear solver, DO:
- 230 3.1.1. Use factorization of compressed matrix and of $K_{\mathcal{S}}$ to solve CP (23a)-(23b) for right-
 231 hand sides $(t_0^\ell, t_{\mathcal{S}}^\ell)$ and RETURN solution $(q_0^\ell, q_{\mathcal{S}}^\ell)$.
- 232 3.1.2. From (17), get ϵ_r^ℓ using solution vector $(\Delta w_0^\ell, \Delta w_{\mathcal{S}}^\ell)$ and right-hand side vectors $(r_0, r_{\mathcal{S}})$.
 233 If $\|\epsilon_r^\ell\| \leq \tau_k$, TERMINATE.
- 234 3.1.3. If $\ell = m_{it}$, increase C , and RETURN to Step 3.1.
- 235 3.2. Recover $(\Delta x_k, \Delta y_k)$ from $(\Delta w_0^\ell, \Delta w_{\mathcal{S}}^\ell)$.
- 236 3.3. Recover $\Delta \nu_k$ from (7).

237 We call our clustering-based IP framework `IP-CLUSTER`. The framework is written in C++
 238 and uses MPI for parallel computations. In this implementation we use the primal-dual IP
 239 algorithm of Mehrotra [19]. We use the matrix templates and direct linear algebra routines
 240 of the `BLOCK-TOOLS` library [15]. This library is specialized to block matrices as those aris-
 241 ing in this work and greatly facilitated the implementation. Within `BLOCK-TOOLS`, we use
 242 its MA57 interface to perform all direct linear algebra operations. We use the GMRES imple-
 243 mentation within the `PETSc` library (<http://www.mcs.anl.gov/petsc>) to perform all it-
 244 erative linear algebra operations. We have implemented serial and parallel versions of CP.
 245 We highlight that the parallel version performs the factorizations of (23b) in parallel and ex-
 246 ploits the block-bordered-diagonal structure of the KKT matrix to perform matrix-vector op-
 247 erations in parallel as well. We use the K-means and hierarchical clustering implementations
 248 of the `C-Clustering` library ([http://bonsai.hgc.jp/~mdehoon/software/cluster/](http://bonsai.hgc.jp/~mdehoon/software/cluster/software.htm)
 249 [software.htm](http://bonsai.hgc.jp/~mdehoon/software/cluster/software.htm)). To implement the market clearing models we use an interface to `AMPL` to cre-
 250 ate individual instances (.nl files) for each scenario and indicate first-stage variables and con-
 251 straints using the `suffix` capability.

252 4.1 Benchmark Problems

We consider stochastic variants of problems obtained from the CUTER library and benchmark problems (SSN, GBD, LANDS, 20TERM) reported in [17]. The deterministic CUTER QP problems have the form

$$\min \frac{1}{2}y^T Qy + d^T y, \text{ s.t. } Ay = b, y \geq 0. \quad (39)$$

We generate a stochastic version of this problem by defining b as a random vector. We create scenarios for this vector $b_s, s \in \mathcal{S}$ using the nominal value b as mean and a standard deviation $\pm\sigma = 0.5b$. We then formulate the two-stage stochastic program:

$$\min e^T y_0 + \sum_{s \in \mathcal{S}} p_s \left(\frac{1}{2}y_s^T Qy_s + d^T y_s \right) \quad (40a)$$

$$\text{s.t. } Ay_s = b_s, s \in \mathcal{S} \quad (40b)$$

$$y_s + y_0 \geq 0, s \in \mathcal{S} \quad (40c)$$

$$y_0 \geq 0. \quad (40d)$$

253 Here, we set $p_s = 1/|\mathcal{S}|$. We first demonstrate the quality of CP in terms of the number of
 254 GMRES iterations. For all cases, we assume a scenario compression rate of 75% (only 25% of the
 255 scenarios are used in the compressed matrix), and we solve the problems to a tolerance of $1 \times$
 256 10^{-6} . We use the notation x% to indicate the compression rate (i.e., the preconditioner CP uses
 257 100-x% of the scenarios to define the compressed matrix). A compression rate of 0% indicates
 258 that the entire scenario set is used for the preconditioner (ideal). A compression rate of 100%
 259 indicates that no preconditioner is used. We set the maximum number of GMRES iterations m_{it}
 260 to 100.

Table 1 Performance of naive and preconditioner CPs in benchmark problems.

Problem	S	n	m	NPI (75%)			CP (75%)			NPII (75%)		
				IPit	LAit	LAit/IPit	IPit	LAit	LAit/IPit	IPit	LAit	LAit/IPit
HS53	100	1,010	800	19	152	8	19	113	5	20	112	5
LOTSCHD	100	1,212	700	27	911*	33	25	203	8	26	178	6
HS76	100	707	300	24	626*	26	23	98	4	24	107	4
HS118	100	5,959	4,400	47	1499*	31	47	409	8	50	484	9
QPCBLEND	100	11,514	7,400	57	258	4	57	253	4	55	273	4
ZECEVIC2	100	606	400	27	451*	16	29	111	3	26	107	4
QPTEST	100	505	300	23	569*	24	23	108	4	26	109	4
SSN	100	70,689	8,600	114	738*	6	114	1857	16	114	2506	21
GBD	1000	10,017	5,000	24	627*	26	24	144	6	24	92	4
LANDS	1000	12,004	7,003	29	481*	16	29	115	3	29	122	4
20TERM	100	76,463	12,404	57	581*	10	57	976	17	58	905*	16

For this first set of results we cluster the scenarios using a hierarchical clustering algorithm with the features (35). The results are presented in Table 1. As can be seen, the performance of CP is satisfactory in all instances, requiring fewer than 20 GMRES iterations per interior-point iteration (this is labeled as LAit/IPit). We attribute this to the particular structure of CP, which enable us to pose the preconditioning systems in the equivalent form (27) and to derive favorable spectral properties and error bounds. To support these observations, we have also

experimented with a couple of *naive* preconditioners. The first naive preconditioner (NPI) has the form:

$$\begin{bmatrix} \bar{K}_S & \bar{B}_S \\ \bar{B}_S^T & K_0 \end{bmatrix} \begin{bmatrix} q_S \\ q_0 \end{bmatrix} = \begin{bmatrix} t_S \\ t_0 \end{bmatrix}, \quad (41)$$

where,

$$\bar{K}_S := \text{blkdiag} \{K_{c_1}, \dots, K_{c_1}, K_{c_2}, \dots, K_{c_2}, \dots, K_{c_C}, \dots, K_{c_C}\} \quad (42a)$$

$$\bar{B}_S := \text{rowstack} \{B_{c_1}, \dots, B_{c_1}, B_{c_2}, \dots, B_{c_2}, \dots, B_{c_C}, \dots, B_{c_C}\}. \quad (42b)$$

We can see that NPI replaces the block matrix elements of the cluster with those of the scenario representing the cluster. This, in fact, is done implicitly by the compressed system (23a). Note also that the right-hand side of NPI is consistent with that of the KKT system. The design of NPI seems reasonable at first sight but it has several structural deficiencies. We highlight these by noticing that the Schur system of NPI has the form

$$\bar{Z}q_0 = t_0 - \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}_i} B_{c_i}^T K_{c_i}^{-1} t_s. \quad (43)$$

This system has the same Schur matrix as that of the Schur system of CP (25) but does not have the same right-hand side. Moreover, the second-stage steps obtained from NPI are

$$K_{c_i} q_s = t_s - B_{c_i} q_0, \quad i \in \mathcal{C}, \quad s \in \mathcal{S}_i. \quad (44)$$

By comparing (44) with (23b) we can see that the recovery of the second-stage steps in NPI does not use the second-stage matrices K_s, B_s corresponding to each scenario (as is done in CP approach). We also consider an alternative naive preconditioner (NPII) to analyze the impact of the second-stage step (23b). This preconditioner computes q_0 using (41) as in NPI but computes the second-stage steps using (44) as in CP. Consequently, NPII and CP only differ in the way q_0 is computed. It is not difficult to verify that the solution of NPII is equivalent to the solution of the system:

$$\begin{bmatrix} K_S & B_S \\ B_S^T & K_0 + E_Z \end{bmatrix} \begin{bmatrix} q_S \\ q_0 \end{bmatrix} = \begin{bmatrix} t_S \\ t_0 + t_{CP} \end{bmatrix}. \quad (45)$$

261 By comparing the equivalent system (27) of CP and the equivalent system (45) of NPII we
262 can see that NPII introduces the additional perturbation t_{CP} on the right-hand side.

263 The structural deficiencies of NPI and NPII prevent us from obtaining the error bounds of
264 Theorems 1 and 2 and highlight the importance of the structure of CP. In Table 1 we compare
265 the performance of the different preconditioners. We can see that the performance of NPI is not
266 competitive. In particular, CP outperforms NPI in nine instances out of eleven. Moreover, in all
267 instances except HS53 and QPCBLEND, it was necessary to refine preconditioner NPI in several
268 iterations (this is done by increasing the number of clusters). We highlight instances in which
269 this occurs using a star next to the total number of GMRES iterations. The performance of NPII
270 is highly competitive with that of CP. In fact, NPII performs slightly better than CP in several
271 instances. For problem 20TERM, however, it was necessary to increase the number of clusters
272 for NPII in some iterations. We can thus conclude that CP has in general better performance and
273 is more robust. Moreover, we can conclude that the second-stage step (23b) plays a key role.

274 In Table 2 we compare the performance of CP with that of the unpreconditioned strategy
275 (compression rate of 100%) and with that of the naive strategies. We only show results for a
276 single instance to illustrate that the matrices of the benchmark problems are nontrivial and
277 preconditioning is indeed needed.

Table 2 Performance of preconditioned and unpreconditioned strategies.

Problem	S	n	Compress.	IPit	LAit	LAit/IPit
HS53	100	1,010	100%	19	12861	676
			75% (NPI)	19	152	8
			75% (CP)	19	113	5
			75% (NPII)	20	112	5

Table 3 Effect of compression rates on 20TERM problem.

S	n	Compress.	IPit	LAit/IPit	θ_{tot}
100	76,463	0%	54	-	16
		50%	57	10	54
		75%	57	19	79
		87%	57	17	69
200	152,863	0%	69	-	44
		50%	72	9	137
		75%	72	14	166
		87%	72	18	185
400	305,663	0%	87	-	108
		50%	92	20	578
		75%	92	21	555
		87%	92	23	570
800	611,263	0%	88	-	232
		50%	97	25	1440
		75%	97	25	1427
		87%	97	27	1417

278 We note that the instances reported in Tables 1 and 2 are small ($n < 100,000$). In most of these
279 small instances we found that the solution times obtained with full factorization are shorter
280 than those obtained with CP. This is because the overhead introduced by the iterative linear
281 solver is not sufficient to overcome the gains obtained by compressing the linear system. We
282 illustrate this behavior in Table 3 where we compare the performance of full factorization (0%
283 compression rate) with that of preconditioner CP for problem 20TERM. We clearly see that the
284 total solution times (denoted as θ_{tot}) obtained with full factorization are significantly shorter
285 than those obtained with CP. Most notably, this trend holds for problems with up to 600,000
286 variables and the times scale linearly with the number of scenarios. These results illustrate that
287 sparse direct factorization codes such as MA57 can efficiently handle certain large-scale prob-
288 lems. As we show in Section 4.2, this efficiency enables us to overcome scalability bottlenecks of
289 Schur decomposition. Full factorization, however, will eventually become expensive as we in-
290 crease the problem size and, at this point, the use of CP becomes beneficial. We illustrate this in
291 Table 6 where we compare the performance of CP with that of full factorization for two large in-
292 stances. Instance QSC2015 has 63,717 variables, while instance AUG3DC has 131,682 variables.
293 We use θ_{fact} to denote the time spent in the factorization of the compressed matrix and of the
294 block matrices, θ_{clus} to denote the time spent performing clustering operations, and θ_{gmres} to
295 denote the time spent in GMRES iterations (without considering factorization operations in the
296 preconditioner). As can be seen, the solution times of full factorization are dramatically reduced
297 by using CP.

298 From Table 3 we can also see that the performance of CP deteriorates as we increase the
299 compression rate. This is because the distortion metric increases as we increase the compres-
300 sion rate and thus the quality of the preconditioner deteriorates, as suggested by Theorems 1
301 and 2. We recall, however, that the bounds provided in these theorems depend on constants

Table 4 Performance of different clustering strategies for benchmark problems (Theorem 1).

Problem	Compress.	c_K	$\sigma_{min}(\bar{Z})$	$\ Z - \bar{Z}\ $	\mathcal{D}_C
LANDS	50%	$6.8 \times 10^{+2}$	3.7×10^{-3}	3.9×10^{-2}	6.5×10^{-2}
	75%	$6.8 \times 10^{+2}$	3.8×10^{-3}	1.8×10^{-1}	5.8×10^{-1}
GBD	50%	$4.5 \times 10^{+12}$	6.8×10^{-2}	6.6×10^{-3}	5.0×10^{-4}
	75%	$4.5 \times 10^{+12}$	6.7×10^{-2}	6.1×10^{-2}	1.6×10^{-3}

Table 5 Performance of different clustering strategies for benchmark problems (Theorem 2).

Problem	Compress.	c_Z	\mathcal{D}_C	$\ (\bar{Z} - Z)Z^{-1}t_Z\ $	$\ (\bar{Z} - Z)t_Z\ $
LANDS	50%	$5.7 \times 10^{+2}$	$1.0 \times 10^{+1}$	$2.0 \times 10^{+0}$	$6.1 \times 10^{+2}$
	75%	$5.8 \times 10^{+2}$	$1.0 \times 10^{+2}$	$2.9 \times 10^{+1}$	$8.5 \times 10^{+3}$
GBD	50%	$9.2 \times 10^{+0}$	1.0×10^{-1}	5.6×10^{-3}	9.2×10^{-2}
	75%	$9.2 \times 10^{+0}$	3.5×10^{-1}	2.0×10^{-2}	7.1×10^{-1}

302 that change with the clustering parameters. Consequently, it is not obvious that reducing the
303 distortion metric will improve the quality of the preconditioner. We designed a numerical ex-
304 periment to gain more insight into this issue. In the experiment we compute the constants and
305 metrics of Theorems 1 and 2 for two additional instances (GBD and LANDS) and for two dif-
306 ferent compression rates (50% and 75%). We only report the results at a single iteration because
307 we observed similar behavior at other iterations. The results are summarized in Tables 4 and 5.
308 As can be seen in Table 4, the constants c_K and $\sigma_{min}(\bar{Z})$ of Theorem 1 are insensitive to the com-
309 pression rate. The distortion metric \mathcal{D}_C , on the other hand, changes by an order of magnitude.
310 We also report the Schur complement error $\|E_Z\| = \|\bar{Z} - Z\|$ and we see that this error changes
311 by an order of magnitude as well. In Table 5 we can see that the constant c_Z of Theorem 2 is
312 insensitive to the compression rate but the distortion is rather sensitive as well. Moreover, we
313 can see that metrics $\|(\bar{Z} - Z)Z^{-1}t_Z\|$, $\|(\bar{Z} - Z)t_Z\|$, and \mathcal{D}_C change significantly with the com-
314 pression rate. We highlight that the distortion metric of Theorem 1 is defined using the features
315 (33) while the distortion metric of Theorem 2 is defined using the features (35). *From these results*
316 *we can conclude that the distortion metrics proposed are indeed appropriate indicators of preconditioning*
317 *quality and can thus be used to guide the construction of the preconditioners.*

318 From Table 3 we can also see that the deterioration of performance due to increasing com-
319 pression rates becomes less pronounced as we increase the number of scenarios. The reason is
320 that more redundancy is observed as we increase the number of scenarios and, consequently,
321 compression potential increases. This behavior has been found in several instances and indi-
322 cates that it is possible to deal with problems with a large number of scenarios.

323 In Table 6 we compare the performance of different clustering strategies. To this end, we per-
324 form clustering using features (33) (we label this as $X^{-1}V$) and using (37) (we label this as r_Z).
325 As can be seen, the performance of both clustering strategies is very similar. This demonstrates
326 that the design of the features (33) and (35) is consistent.

327 4.2 Stochastic Market Clearing Problem

We demonstrate the computational efficiency of the preconditioner by solving a stochastic market-
clearing model for the entire Illinois power grid system [21,25,27]. The system is illustrated in

Table 6 Performance of different clustering strategies for benchmark problems.

Problem	Compress.	Clustering	IPit	θ_{tot}	θ_{fact}	θ_{clus}	θ_{gmres}	LAit	LAit/IPit
QSC205	0%		110	1331	1321				
	50%	$X^{-1}V$	110	220	157	5	42	747	6
	75%	$X^{-1}V$	110	91	25	6	45	933	8
	50%	r_Z	110	229	161	5	43	747	6
	75%	r_Z	110	89	24	5	43	924	8
AUG3DC	0%		11	1427	1423				
	50%	$X^{-1}V$	11	96	84	0.3	6	26	2
	75%	$X^{-1}V$	11	24	13	0.3	6	27	2
	50%	r_Z	11	93	80	0.3	6	26	2
	75%	r_Z	11	25	13	0.3	6	27	2

Figure 1. The stochastic programming formulation is given by

$$\begin{aligned} \min_{x_i, X_i(s)} \sum_{i \in \mathcal{G}} \left(\beta_i x_i + \sum_{s \in \mathcal{S}} p_s [\beta_i^+ (X_i(s) - x_i)_+ - \beta_i^- (X_i(s) - x_i)_-] \right) \\ \text{s.t. } \tau_n(f) + \sum_{i \in \mathcal{G}(n)} x_i = d_n, \quad n \in \mathcal{N} \end{aligned} \quad (46a)$$

$$\tau_n(F(s)) - \tau_n(f) + \sum_{i \in \mathcal{G}(n)} (X_i(s) - x_i) = 0, \quad n \in \mathcal{N}, s \in \mathcal{S} \quad (46b)$$

$$f, F(s) \in \mathcal{F}, \quad s \in \mathcal{S} \quad (46c)$$

$$(x_i, X_i(s)) \in \mathcal{X}_i(s), \quad i \in \mathcal{G}, s \in \mathcal{S}. \quad (46d)$$

Here, \mathcal{N} denotes the set of network nodes and the set of all suppliers is denoted by \mathcal{G} . Subsets $\mathcal{G}(n)$ denote the set of suppliers connected to node $n \in \mathcal{N}$. The forward (first-stage) dispatched quantities for players are x_i , and the spot (second-stage) quantities under scenario ω are $X_i(s)$. Symbol f denotes the vector of all line flows and $\tau_n(\cdot)$ are flow injections into node $n \in \mathcal{N}$. Similarly, $F(s)$ denotes the vector of line flows for each scenario s . The demand is assumed to be deterministic and inelastic and is represented by d_n , $n \in \mathcal{N}$. The sets \mathcal{F} and $\mathcal{X}_i(s)$ are polyhedral and define lower and upper bounds for the flows and dispatch quantities. The objective of the market clearing problem is to minimize the forward dispatch cost plus the expected recourse dispatch cost. Here $[y]_+ = \max\{y, 0\}$ and $[y]_- = \max\{-y, 0\}$. The coefficients β_i denote the supply price bids, and β_i^+ and β_i^- are price bids for corrections of the suppliers. A supplier i asks $\beta_i^+ > \beta_i$ to sell additional power or asks $\beta_i^- < \beta_i$ to buy power from the system (e.g., reduces output). The scenarios $s \in \mathcal{S}$ characterize the randomness in the model due to unpredictable supply capacities (in this case wind power).

The market clearing model has *large first-stage dimensionality*. The Schur complement has a dimension of 64,199 and has a large dense block. In Table 7 we present the solution times for this problem using a Schur decomposition strategy for *a single scenario*. Here, θ_{tot} is the total solution time, $\theta_{factschur}$ is the time spent factorizing the Schur complement, $\theta_{formschur}$ is the time spent forming the Schur complement, and $\theta_{factblock}$ is the time spent factorizing the scenario blocks (in this case just one block). All times are reported in seconds. The solution time for this problem is 3.61 hr, with 25% of the time spent forming the Schur complement and 75% spent factorizing the Schur complement. Note that if more scenarios are added, the time spent forming and factorizing the Schur complement will dominate (even if the scenarios can be parallelized). Iterative strategies applied to the Schur complement system can avoid the time

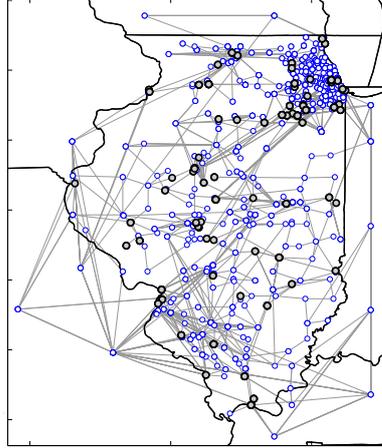


Fig. 1 Illinois transmission network. Dark dots are supply nodes and blue dots are demand nodes.

Table 7 Performance of Schur decomposition approach.

S	n	IPit	θ_{tot}	$\theta_{factschur}$	$\theta_{formschur}$	$\theta_{factblock}$
1	30,472	55	31280	27236	4023	4

351 spent forming the Schur complement but not the factorization time because a preconditioner
 352 with a large dense block still needs to be factored [20].

353 We now assess the serial performance of CP. By comparing Tables 7 and 8 we can see that
 354 the full factorization approach will be as efficient as Schur decomposition for problems with up
 355 to 64 scenarios. In other words, it would be faster to factorize the full sparse KKT system than
 356 forming and factorizing the large Schur complement. The fast growth in solution time of the
 357 full factorization approach is remarkable, however. We attribute this to the tight connectivity
 358 induced by the network constraints which introduce significant fill-in. CP reduces the solution
 359 times of full factorization by a factor of 2 for the problem with 32 scenarios and by a factor of 8
 360 for the problem with 64 scenarios. We highlight that CP is highly effective, requiring on average
 361 5 to 10 GMRES iterations per IP iteration for compression rates of 50% and 11 to 13 iterations for
 362 compression rates of 75%. We also observe that the performance of different clustering strategies
 363 is similar.

364 For the problem with 64 scenarios we can see that the solution time of CP increases as we
 365 increase the compression rate from 75% to 87% (even if the factorization time is dramatically
 366 reduced). This is because the time spent in GMRES to perform backsolves and matrix-vector
 367 operations dominates the factorization time. We mitigate this by using the parallel implementa-
 368 tion of CP. The results are presented in Table 9. We can see that the solution time spent in GMRES
 369 to perform backsolves and matrix-vector operations is dramatically reduced by exploiting the
 370 block-bordered-diagonal structure of the KKT matrix. *This enables us to solve a market clearing*
 371 *problem with over 1.2 million variables in 10 minutes, as opposed to 9 hours using the full factorization*
 372 *approach. This represents a speed up factor of 42.* By comparing the parallel results with those of
 373 Table 7 we can also see that the Schur complement approach is not competitive because of the
 374 time needed to form and factorize the Schur complement (this holds even for a single scenario).

375 From Table 9 we can see that scalability slows down as we increase the number of processes.
 376 This is because the remaining serial components (beyond backsolves and matrix-vector opera-

Table 8 Serial performance of preconditioner CP against full factorization for stochastic market clearing problem.

S	n	Compress.	Cluster.	IPit	θ_{tot}	θ_{fact}	θ_{clus}	θ_{gmres}	LAit	LAit/IPit
16	309,187	0%		57	473	452				
		50%	$X^{-1}V$	57	544	119	0.4	325	631	11
		50%	r_Z	57	508	117	0.15	296	519	9
32	606,483	0%		65	3480	3414				
		50%	$X^{-1}V$	65	1477	661	8	606	574	8
		75%	$X^{-1}V$	65	1479	145	8	1141	1194	18
		50%	r_Z	65	1347	672	3	459	398	6
		75%	r_Z	65	1131	150	3	769	804	12
64	1,201,075	0%		64	28022	27883				
		50%	$X^{-1}V$	64	5163	3513	29	1292	660	10
		75%	$X^{-1}V$	64	2878	656	29	1844	902	14
		87%	$X^{-1}V$	64	2499	135	29	1990	1040	16
		50%	r_Z	64	5238	3492	12	1349	666	10
		75%	r_Z	64	3003	659	12	1924	937	14
		87%	r_Z	64	2440	115	12	1944	1147	17

Table 9 Parallel performance of preconditioner CP against full factorization for stochastic market clearing problem.

S	n	MPI Proc.	Compress.	IPit	θ_{tot}	θ_{fact}	θ_{clus}	$\theta_{factblock}$	θ_{gmres}	LAit	LAit/IPit
64	1,201,075	1	0%	64	28022	27883					
		1	87%	64	2440	115	12	288	1944	1147	17
		2	87%	64	1211	116	12	147	892	1025	16
		4	87%	64	817	134	12	80	592	919	14
		8	87%	64	658	152	12	44	398	905	14
		1	94%	64	3223	49	12	327	2764	1489	23
		2	94%	64	1558	43	12	151	1306	1471	22
		4	94%	64	993	49	12	84	801	1420	22
		8	94%	64	733	54	12	46	570	1409	22

377 tions) of CP start dominating. This overhead includes operations inside the GMRES algorithm
378 itself. We are currently investigating ways to parallelize these operations.

379 In Table 9 we also present experiments using a compression rate of 94%. We performed
380 these experiments to explore the performance limit of CP. We can see that the performance of
381 CP deteriorates in terms of total solution time because the number of GMRES iterations (and
382 thus time) increases. Consequently, it does not pay off to cluster the KKT system further. We
383 highlight, however, that the deterioration of CP in terms of GMRES iterations is *graceful*. It is
384 remarkable that, on average, the linear system can be solved in 22 GMRES iterations when
385 only four scenarios are used in the compressed matrix. This behavior again indicates that the
386 computation of the second-stage variables in (23b) plays a key role in the performance of CP.

387 We emphasize on the efficiency gains obtained from parallelization with respect to the com-
388 putation of the second-stage steps (23b). This step requires a factorization of all the block ma-
389 trices K_s prior to calling the iterative linear solver. When the factorizations of the blocks are
390 performed serially, the total solution time grows linearly with the number of scenarios. This
391 can be observed from the block factorization times (denoted as $\theta_{factblock}$) reported in Table 9. In
392 particular, the time spent in the factorization of the block matrices in the serial implementation
393 (one processor) is a significant component of the total time. This overhead is eliminated using
394 the parallel implementation (with almost perfect scaling).

395 5 Conclusions and Future Work

396 We have presented a preconditioning strategy for stochastic programs using clustering tech-
 397 niques. This inside-the-solver clustering strategy can be used as an alternative to (or in combi-
 398 nation with) outside-the-solver scenario aggregation and clustering strategies. Practical features
 399 of performing inside-the-solver clustering is that no information on probability distributions is
 400 required and the effect of the data on the problem at hand is better captured. We have demon-
 401 strated that the preconditioners can be implemented in sparse form and dramatically reduce
 402 computational time compared to full factorizations of the KKT system. We have also demon-
 403 strated that the sparse form enables the solution of problems with large first-stage dimension-
 404 ality that cannot be addressed with Schur decomposition. Scenario compression rates of up to
 405 87% have been observed in large problem instances. As part of future work, we will investi-
 406 gate the performance of the preconditioner in a nonlinear programming setting and we will
 407 investigate extensions to multi-stage stochastic programs.

408 **Acknowledgements** This material was based upon work supported by the U.S. Department of Energy, Office of Sci-
 409 ence, under Contract No. DE-AC02-06CH11357. Victor M. Zavala acknowledges funding from the DOE Office of Science
 410 under the Early Career program. Carl Laird and Yankai Cao acknowledge support by the National Science Foundation
 411 CAREER Grant CBET #0955205. The authors thank Jacek Gondzio for providing feedback on a previous version of the
 412 manuscript.

413 References

- 414 1. J. Birge. Aggregation bounds in stochastic linear programming. *Mathematical Programming*, 31:25–41, 1985.
- 415 2. C.M. Bishop et al. *Pattern recognition and machine learning*, volume 4. Springer, New York, 2006.
- 416 3. R. H. Byrd, G. M. Chin, W. Neveitt, and J. Nocedal. On the use of stochastic Hessian information in optimization
 417 methods for machine learning. *SIAM Journal on Optimization*, 21(3):977–995, 2011.
- 418 4. G. C. Calafiore and M. C. Campi. The scenario approach to robust control design. *Automatic Control, IEEE Transac-*
 419 *tions on*, 51(5):742–753, 2006.
- 420 5. M.S. Casey and S. Sen. The scenario generation algorithm for multistage stochastic linear programming. *Mathemat-*
 421 *ics of Operations Research*, 30(3):615–631, 2005.
- 422 6. Naiyuan Chiang and Andreas Grothey. Solving security constrained optimal power flow problems by a structure
 423 exploiting interior point method. *Optimization and Engineering*, pages 1–23, 2012.
- 424 7. M. Colombo, J. Gondzio, and A. Grothey. A warm-start approach for large-scale stochastic linear programs. *Math-*
 425 *ematical Programming*, 127(2):371–397, 2011.
- 426 8. W. L. de Oliveira, C. Sagastizábal, D. Penna, M. Maceira, and J. M. Damázio. Optimal scenario tree reduction for
 427 stochastic streamflows in power generation planning problems. *Optimization Methods and Software*, 25(6):917–936,
 428 2010.
- 429 9. HS Dollar. Constraint-style preconditioners for regularized saddle point problems. *SIAM Journal on Matrix Analysis*
 430 *and Applications*, 29(2):672–684, 2007.
- 431 10. J. Dupačová, N. Gröwe-Kuska, and W. Römisch. Scenario reduction in stochastic programming. *Mathematical*
 432 *programming*, 95(3):493–511, 2003.
- 433 11. M. C. Ferris and T. S. Munson. Interior-point methods for massive support vector machines. *SIAM Journal on*
 434 *Optimization*, 13(3):783–804, 2002.
- 435 12. J. Gondzio and A. Grothey. Reoptimization with the primal-dual interior point method. *SIAM Journal on Optimiza-*
 436 *tion*, 13:842–864, 2003.
- 437 13. H. Heitsch and W. Römisch. Scenario tree reduction for multistage stochastic programs. *Computational Management*
 438 *Science*, 6:117–133, 2009.
- 439 14. J. Jung, D. P. O’Leary, and A. L. Tits. Adaptive constraint reduction for training support vector machines. *Electronic*
 440 *Transactions on Numerical Analysis*, 31:156–177, 2008.
- 441 15. J. Kang, Y. Cao, D. P. Word, and C.D. Laird. An interior-point method for efficient solution of block-structured NLP
 442 problems using an implicit Schur-complement decomposition. *Computers & Chemical Engineering*, In Press, 2014.
- 443 16. J. M. Latorre, S. Cerisola, and A. Ramos. Clustering algorithms for scenario tree generation: Application to natural
 444 hydro inflows. *European Journal of Operational Research*, 181(3):1339 – 1353, 2007.
- 445 17. J. Linderoth, A. Shapiro, and S. Wright. The empirical behavior of sampling methods for stochastic programming.
 446 *Annals of Operations Research*, 142(1):215–241, 2006.

- 447 18. M. Lubin, C. G. Petra, M. Anitescu, and V. M. Zavala. Scalable stochastic optimization of complex energy systems.
448 In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, pages 1–10. IEEE,
449 2011.
- 450 19. S. Mehrotra. On the implementation of a primal-dual interior point method. *SIAM Journal on Optimization*, 2:575–
451 601, 1992.
- 452 20. C. Petra and M. Anitescu. A preconditioning technique for Schur complement systems arising in stochastic opti-
453 mization. *Computational Optimization and Applications*, 52:315–344, 2012.
- 454 21. G. Pritchard, G. Zakeri, and A. Philpott. A single-settlement, energy-only electric power market for unpredictable
455 and intermittent participants. *Operations Research*, 58(4-part-2):1210–1219, 2010.
- 456 22. C. M. Shetty and R. W. Taylor. Solving large-scale linear programs by aggregation. *Computers & Operations Research*,
457 14(5):385 – 393, 1987.
- 458 23. D. B. Szyld and J. A. Vogel. Fqmr: A flexible quasi-minimal residual method with inexact preconditioning. *SIAM*
459 *Journal on Scientific Computing*, 23(2):363–380, 2001.
- 460 24. A. Tits, P. Absil, and W. Woessner. Constraint reduction for linear programs with many inequality constraints.
461 *SIAM Journal on Optimization*, 17(1):119–146, 2006.
- 462 25. V. M. Zavala, A. Botterud, E. M. Constantinescu, and J. Wang. Computational and economic limitations of dispatch
463 operations in the next-generation power grid. *IEEE Conference on Innovative Technologies for and Efficient and Reliable*
464 *Power Supply*, 2010.
- 465 26. V. M. Zavala, E. M. Constantinescu, T. Krause, and M. Anitescu. On-line economic optimization of energy systems
466 using weather forecast information. *Journal of Process Control*, 19(10):1725–1736, 2009.
- 467 27. V. M. Zavala, K. Kim, M. Anitescu, and J. Birge. A stochastic market clearing formulation with consistent pricing
468 properties. *Technical Report ANL/MCS-P5110-0314, Argonne National Laboratory*, 2015.
- 469 28. P.H. Zipkin. Bounds for row-aggregation in linear programming. *Operations Research*, 28(4):903–916, 1980.