

Noname manuscript No.
(will be inserted by the editor)

A preconditioning technique for Schur complement systems arising in stochastic optimization

Cosmin G. Petra · Mihai Anitescu

Received: date / Accepted: date

Abstract Deterministic sample average approximations of stochastic programming problems with recourse are suitable for a scenario-based, treelike parallelization with interior-point methods and a Schur complement mechanism. However, the direct linear solves involving the Schur complement matrix are expensive, and adversely affect the scalability of this approach. In this paper we propose a stochastic preconditioner to address this issue. The spectral analysis of the preconditioned matrix indicates an exponential clustering of the eigenvalues around 1. The numerical experiments performed on the relaxation of a unit commitment problem show good performance, in terms of both the accuracy of the solution and the execution time.

Keywords stochastic programming · saddle-point preconditioning · Krylov methods · interior-point method · sample average approximations · parallel computing

1 Introduction

Stochastic programming (SP) is concerned with solving optimization problems involving uncertainty in both the objective and the constraints. In this paper we consider two-stage stochastic convex quadratic problems with recourse. These are mathematically formulated as

$$\min \left(\frac{1}{2} x_0^T Q_0 x_0 + c_0^T x_0 \right) + \mathcal{G}(x_0) \quad \text{subject to} \quad T_0 x_0 = b_0, \quad x_0 \geq 0, \quad (1)$$

where $\mathcal{G}(x_0) = \mathbb{E}[G(x_0, \xi(\omega))]$ and $G(x_0, \xi(\omega))$ is the optimal value of the second-stage problem

$$\min \frac{1}{2} y^T Q(\omega) y + c(\omega)^T y \quad \text{subject to} \quad W(\omega) y = b(\omega) - T(\omega) x_0, \quad y \geq 0. \quad (2)$$

C. Petra · M. Anitescu
Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Avenue, Argonne, IL 60439, USA
E-mail: petra,anitescu@mcs.anl.gov

Here $\xi(\omega) := (Q(\omega), c(\omega), T(\omega), W(\omega), b(\omega))$ denotes the data of the second-stage problem and is viewed as a random vector defined over a probability space Ω with a known probability distribution P . We also assume that $Q(\omega)$ is symmetric positive semidefinite and that both the technology matrix $T(\omega)$ and the recourse matrix $W(\omega)$ have full rank for any ω .

The first stage of the problem is the present moment of time, while the second stage corresponds to a future point in time at which the uncertainty is revealed. However, a decision x_0 must be taken before the actual realizations of the uncertain parameters of the second stage become available. The optimal decision x_0 also must take into account a “recourse” action for the possible outcome of the uncertainty. In the context of stochastic programming problems with recourse, the objective consists of not only the usual “operation cost” $\left(\frac{1}{2}x_0^T Q_0 x_0 + c_0^T x_0\right)$ but also the expected value $\mathbb{E}[G(x_0, \xi(\omega))]$ of the “recourse cost” taken with respect to all possible outcomes of the uncertainty.

In the case of a finitely supported distribution P , that is, there exist a finite number of realizations $\xi_1, \xi_2, \dots, \xi_N$ of the random vector $\xi(\omega)$, the recourse function is a weighted average

$$\mathcal{G}(x_0) = \sum_{i=1}^N p_i G(x_0, \xi_i),$$

where p_i is the probability of ξ_i to occur, $i = 1, 2, \dots, N$. Consequently, the SP problem (1) is a deterministic optimization problem with a special structure, but standard derivative-based cannot be applied because of the nonsmooth (piecewise differentiable, more exactly) recourse term. The reader may refer to [3, 6] for a detailed survey of algorithms for SP problems. The L-shaped method [43] and its variants use Benders’ decomposition of the primal problem or a Dantzig-Wolfe decomposition of the dual to exploit the special structure of the two-stage SP problem. Based on cutting-plane subgradient approximation of the objective, this class of methods builds a lower convex, nondifferentiable estimate of the recourse cost function and iterate by minimizing this estimate. Stabilization techniques such as quadratic regularization [41] or an l_∞ trust-region approach [27] are needed to make this class of methods robust. Among the alternatives to subgradient-based methods we mention Lagrangian-based approaches [38, 26, 39, 40] and direct approaches that use simplex or an interior-point method to solve the equivalent deterministic convex quadratic programming problem

$$\begin{aligned} \min \quad & \left(\frac{1}{2}x_0^T Q_0 x_0 + c_0^T x_0\right) + \sum_{i=1}^N p_i \left(\frac{1}{2}x_i^T Q_i x_i + c_i^T x_i\right) \\ \text{subj to} \quad & T_0 x_0 = b_0, \\ & T_i x_0 + W_i x_i = b_i, \quad i = 1, 2, \dots, N, \\ & x_i \geq 0, \quad i = 0, 1, \dots, N. \end{aligned} \tag{3}$$

The direct approach is competitive with other methods only if the optimization method takes into account the special form of the problem (3). Generally, each approach has advantages and drawbacks. In this paper we use a primal-dual interior-point method (IPM) for solving problems of form (3). This class of methods provides a unified framework for linear, quadratic and nonlinear continuous programming problems and requires virtually the same linear algebra. Based on Newton’s method, IPMs exhibit polynomial complexity for convex programming problems and local superlinear or quadratic convergence for a large class of problems. They also have emerged in the past twenty-five years as a practical method for solving large-scale optimization problems.

For SP problems having a continuous probability distribution, the evaluation of the recourse function $\mathcal{G}(x_0)$ requires integration of the nonsmooth function $G(x_0, \xi(\omega))$ in the space of the random parameters $\xi(\omega)$ [6]. Even if the integrand $G(x_0, \xi(\omega))$ is easily evaluated, the integration is usually untractable from a computational point of view even for a relatively small size of $\xi(\omega)$. On the other hand, when the probability space $\Omega = \{\omega_1, \omega_2, \dots\}$ is discrete, the expectation term $\mathcal{G}(x_0)$ is a weighted average of $G(x_0, \xi(\omega_i))$, $i = 1, 2, \dots$, and therefore the second stage has to be solved for each possible outcome of Ω , possibly infinitely many times if Ω is not finite. Even in the finite case, the cardinality of Ω can be extremely large since it increases exponentially with the number of independent random variables of $\xi(\omega)$. In order to deal with the issue of numerical integration for the continuous case and infinite or extremely large cardinality in the discrete case, sampling methods are used to obtain a discrete finite approximation of the randomness that is tractable from a computational point of view. Among the sampling techniques used in the context of stochastic programming, we mention Monte Carlo methods (covered in [6, Chapter 10] and [42, Chapter 5]), the Latin hypercube sampling method [30], and the importance sampling method [10].

Two fundamental approaches emerge from the interplay of optimization and sampling. The first is the “interior sampling” approach that performs the optimization directly on (1) and relies on sampling to approximate the recourse function and its derivative information. Methods following this approach include the stochastic quasi-gradient method [13], the stochastic Newton method [4], the stochastic decomposition method [24], and the L-shaped method based on importance sampling of [10]. Unfortunately, repeated samplings make this approach infeasible for applications in which the cost of the sampling is comparable with or larger than the cost of optimization. Weather-based simulation and control problems [8] are such examples.

In the second approach, “exterior sampling”, the sampling is done only once at the beginning of the optimization process. A finite sample $\xi_1, \xi_2, \dots, \xi_N$ of N realizations (scenarios) of the random vector $\xi(\omega)$ is used to approximate the expected value $\mathbb{E}[G(x, \xi(\omega))]$ by the averaging values $G(x_0, \xi_i)$, $i = 1, 2, \dots, N$. The so-called sample average approximation (SAA) problem obtained with this approach is a convex quadratic deterministic programming problem having the form (3) of a discrete SP problem with equal probabilities $p_1 = \dots = p_N = 1/N$. From a computational point of view, the SAA problem is a stochastic programming problem with a finite number of realizations. Therefore, any of the abovementioned numerical methods for the discrete SP problems can be applied to SAA SP problems.

SAA problems can become extremely large. For example in [17], an SAA problem having more than 1 billion variables was solved by using a primal-dual interior-point method. SAA SP problems of this size and even considerably smaller sizes can be handled only in a parallel environment with distributed memory, however, this requires the problem to be decomposed into subproblems that can be solved independently on different processors. In the context of subgradient-based or Lagrangian-based methods, the decomposition is obtained by evaluating the recourse function and obtaining derivative information independently for each scenario. Once an approximation of the recourse term is available, the optimization in the space of x_0 is done in serial and thus causes a bottleneck in the parallel execution flow. On the other hand, the decomposition of the problem in the context of IPMs is usually achieved at the linear algebra level by taking advantage of the block-separable form of the objective function and the half-arrow shape of the Jacobian. This special structure allows most of the work related to IPM linear solves to be done independently for each scenario when a

Schur complement mechanism is used. Decompositions of the SP problems based on the Schur complement and interior-point methods can be found as early as 1988 in the work of Birge and Qi [7]. Recently, IPM-based decomposition for SP problems was implemented in state-of-the-art software packages such as OOPS [18–20] and IPOPT [46]. A slightly different IPM decomposition is the log-barrier Benders-like decomposition introduced by Zhao in [49] for linear problems and extended by Mehrotra and Ozevin to the quadratic case [33] and to two-stage stochastic conic problems [32,34]. A potential advantage of the log-barrier Benders decomposition consists of adaptively adding or removing scenarios during the optimization. However, the two IPM approaches share the same linear algebra, and the technique presented in this paper applies equally to both.

In the early 1990s, because of the limited network bandwidth, the IPM decomposition based on the Schur complement did not prove to be very scalable since it moves large matrices across processors. With the new high-speed interconnects, however, the network bandwidth does not have the same limiting effect, and the main drawback of IPM decomposition is similar to the one of L-shaped methods, namely, the linear algebra of the first stage cannot be started until all the second-stage information is available. More specifically, the only work that cannot be done independently for each scenario consists of linear solves with the Schur complement of the first-stage Hessian subblock in the entire Hessian matrix. The Schur complement for two-stage stochastic programming is a saddle-point matrix of the form

$$\begin{bmatrix} H & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} r_1 \\ r_2 \end{bmatrix},$$

where the block $H \in \mathbb{R}^{n \times n}$ is almost completely dense (n is the number of the first-stage variables) and A represents the Jacobian of the first-stage equalities constraints. The presence of the dense block H incurs a cost of approximately $O(n^3)$ for each linear solve with the Schur complement matrix. For problems having a large number of first-stage variables, the $O(n^3)$ cost becomes dominant for a relatively large number of processors and adversely affects the scalability of the IPM decomposition method.

In this paper we propose an alternative method for the solution of the Schur complement linear systems. The goal of the current work is to remove the $O(n^3)$ cost of the direct factorization of the Schur matrix associated with the first-stage variables x_0 from the parallel execution flow. We substitute the direct solve of the Schur complement matrix with Krylov subspace iterative methods. In the context of the ill-conditioned saddle-point linear systems arising from the primal-dual interior-point algorithms, Krylov methods are known to perform well only when used with a preconditioner. A popular preconditioning approach for the saddle-point linear systems is the constraint preconditioner [25,28,36] that replaces the (1,1) block of the matrix by an easily invertible approximation matrix, while keeping the constraints blocks (1,2) and (2,1) intact. For example, diagonal approximations of the (1,1) block have been used in [2], and implicit factorizations of the same block have been proposed in [11,12]. However these preconditioners are not suitable for the parallel computation framework of this paper because they require data from all the scenarios. Therefore, the factorization of the preconditioner can be started only after all scenarios are finished and hence creates a bottleneck.

The *stochastic preconditioner* we propose in this paper approximates the (1,1) block from the Schur complement matrix by incorporating only a subset of available

scenarios of the SAA problem. Factorization of the preconditioner is carried out on a separate process *before all scenarios are finished* and the bottleneck is thus removed. We analyze the preconditioner from the perspective of the theory of large deviations and show that the eigenvalues of the preconditioned matrix cluster exponentially fast around 1 with the number of the scenarios incorporated in the preconditioner. In practice, the preconditioner allows the Krylov methods to solve the ill-conditioned IPM linear systems with the same accuracy as a direct solver in a small number of iterations.

The paper is organized as follows. In Section 2 we present the linear algebra required by interior-point methods for the solution of the SAA problem. The stochastic preconditioners are introduced and analyzed in Section 3, and the details of their implementation are presented in Section 4. In Section 5 we investigate and report on the practical performance of the preconditioners. The conclusions of this work and related future research directions are given in Section 6.

2 Linear algebra in interior-point methods

In this section we present the linear algebra operations needed by path-following interior-points methods to solve convex quadratic programming problems of the form

$$\min \frac{1}{2}x^T Qx + c^T x \text{ subject to } Ax = b, x \geq 0. \quad (4)$$

In what follows, we suppose that $c, x \in \mathbb{R}^n$, $Q \in \mathbb{R}^{n \times n}$ is symmetric positive semidefinite, $A \in \mathbb{R}^{m \times n}$, and $b \in \mathbb{R}^m$.

Path-following interior-point methods make use of the “central path”, which is a continuous perturbation of the KKT system corresponding to (4) parameterized by $\mu > 0$,

$$\begin{aligned} Qx - A^T y - u &= -c, \\ Ax &= b, \\ xu &= \mu e, \\ x, u &> 0. \end{aligned} \quad (5)$$

Here $y \in \mathbb{R}^m$ and $u \in \mathbb{R}^n$ correspond to the Lagrange multipliers, $e = [1 \ 1 \ \dots \ 1]^T \in \mathbb{R}^n$ and xu denotes the componentwise product.

In the case of a feasible problem (4), the above system has a unique solution (x_μ, y_μ, u_μ) for any $\mu > 0$, and as μ approaches zero, (x_μ, y_μ, u_μ) approaches a maximal complementarity solution of (4) [44]. A path-following method is an iterative numerical process that follows the central path in the direction of decreasing μ toward the solution set of the problem. The iterates generated by the method do not necessarily stay on the central path. Rather, they are located in a controlled neighborhood of the central path that is a subset of the positive orthant.

In the past two decades, predictor-corrector methods have emerged as practical path-following IPMs in solving linear and quadratic programming problems. Among the predictor-corrector methods, the most successful is Mehrotra’s predictor-corrector algorithm. Although Mehrotra [31] presented his algorithm in the context of linear programming, it has been successfully applied also to convex quadratic programming [14] and standard monotone linear complementarity problems [47]. It also has been widely used in the implementation of several IPM-based optimization packages, including

OB1 [29], HOPDM [16], PCx [9], LIPSOL [48], and OOQP [14]. We now describe the Mehrotra's algorithm implemented in OOQP. In pursuing the central path, Mehrotra's predictor-corrector method steps along two directions. The first is the affine-scaling direction, which aims for a great amount of progress toward the solution and generates a point situated on the boundary of the neighborhood. At the k th iteration (x_k, y_k, u_k) of the IPM, the predictor direction $(\Delta x_k^a, \Delta y_k^a, \Delta u_k^a)$ is the Newton direction applied to the nonlinear system (5) with $\mu = 0$, that is,

$$\begin{aligned} Q\Delta x_k^a - A^T \Delta y_k^a - \Delta u_k^a &= -c - Qx_k + A^T y_k + u_k \\ A\Delta x_k^a &= b - Ax_k \\ U_k \Delta x_k^a + X_k \Delta u_k^a &= -x_k u_k. \end{aligned} \quad (6)$$

Here X_k and U_k denote the diagonal matrices with the (positive) entries given by x_k and u_k , respectively. The maximum step-length α is computed such that the predictor update

$$(x_k^a, y_k^a, u_k^a) := (x_k + \alpha \Delta x_k^a, y_k + \alpha \Delta y_k^a, u_k + \alpha \Delta u_k^a)$$

remains in the positive orthant. Mehrotra's algorithm does not update the iterate but computes the centering parameter $\sigma = (\mu_k^a/\mu_k)^3$, where $\mu_k^a = ((x_k^a)^T u^a)/n$.

The corrector direction is biased toward the central path in an attempt to bring the iterate close to central path. It comprises a centering component ($\mu = (\frac{\mu_k^a}{\mu_k})^3$ in (5)), an error-corrector component ($-\Delta x_k^a \Delta u_k^a y_k$ in the right hand side of the third equation from (5)), and the predictor direction itself. More exactly, the corrector direction is the solution of the linear system

$$\begin{aligned} Q\Delta x_k - A^T \Delta y_k - \Delta u_k &= -c - Qx_k + A^T y_k + u_k \\ A\Delta x_k &= b - Ax_k \\ U_k \Delta x_k + X_k \Delta u_k &= \sigma e - \Delta x_k^a \Delta u_k^a y_k - x_k u_k. \end{aligned} \quad (7)$$

The step-length α_k is computed as in the predictor phase and is reduced by a fixed factor in order to keep the new iterate $x_{k+1} = x_k + \alpha_k \Delta x_k$, $y_{k+1} = y_k + \alpha_k \Delta y_k$, and $u_{k+1} = u_k + \alpha_k \Delta u_k$ away from the boundary of the positive orthant.

By performing block elimination for Δu_k , the linear systems (6) and (7) can be reduced to the following symmetric indefinite linear system

$$\begin{bmatrix} Q + D_k^2 & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta x_k \\ -\Delta y_k \end{bmatrix} = \begin{bmatrix} r_k^1 + X_k^{-1} r_k^3 \\ r_k^2 \end{bmatrix}, \quad (8)$$

where $D_k = X_k^{-\frac{1}{2}} U_k^{\frac{1}{2}}$ and r_k^1, r_k^2 , and r_k^3 are the right-hand sides of (6) or (7).

The linear system (7) is known as the "augmented system" and is typically solved by performing a Bunch-Parlett factorization. An alternative is the so-called "normal equations" approach which reduces the augmented system to a symmetric positive definite system of the form $A(Q + D_k)^{-1} A^T \Delta y_k = r$ and employs a Cholesky-based factorization coupled to a mechanism that uses Sherman-Morison-Woodbury formula to deal with the dense columns of A . The normal equations approach is very popular in the context of linear programming ($Q = 0$), but rarely used for quadratic programming because the factorization of $Q + D_k$ and the multiple solves with its factors tend to be computationally burdensome.

2.1 The Schur complement decomposition

We now show how the parallelization of linear algebra can be achieved for two-stage SAA SP problems via the Schur complement mechanism. The technique parallelizes the solution of the augmented system. For similar approaches for the solutions of normal equations for linear programming we refer the reader to [7, 5, 37].

The deterministic problem (3) corresponding to the general two-stage stochastic programming problem given by (1) and (2) has a staircase structure that can be exploited to produce highly parallelizable linear algebra. More specifically, (3) is in fact

$$\begin{aligned}
& \min \left(\frac{1}{2} x_0^T Q_0 x_0 + c_0^T x_0 \right) + \frac{1}{N} \sum_{i=1}^N \left(\frac{1}{2} x_i^T Q_i x_i + c_i^T x_i \right) \\
& \text{subj to } \begin{array}{rcl} T_0 x_0 & & = b_0, \\ T_1 x_0 + W_1 x_1 & & = b_1, \\ T_2 x_0 + & W_2 x_2 & = b_2, \\ \vdots & \ddots & \vdots \\ T_N x_0 + & & W_N x_N = b_N, \\ x_0 \geq 0, x_1 \geq 0, x_2 \geq 0, \dots, x_N \geq 0. \end{array} \quad (9)
\end{aligned}$$

The central path (5) has the following form in the case of the SAA SP problem:

$$\begin{aligned}
Q_0 x_0 - T_0^T y_0 - \frac{1}{N} \sum_{i=1}^N T_i^T y_i - u_0 &= -c_0, \\
\frac{1}{N} Q_i x_i - \frac{1}{N} W_i^T y_i - \frac{1}{N} u_i &= -\frac{1}{N} c_i, \\
T_0 x_0 &= b_0, \\
T_i x_0 + W_i x_i &= b_i, \quad i = 1, 2, \dots, N, \\
x_i u_i &= \mu e, \quad i = 0, 1, \dots, N, \\
x_i, u_i &\geq 0, \quad i = 0, 1, \dots, N.
\end{aligned} \quad (10)$$

In order to be consistent with the optimality conditions for stochastic programming problems presented in [42], the multipliers y_1, y_2, \dots, y_N and u_1, u_2, \dots, u_N were scaled by N in the derivation of (10).

The linear system (8) solved at each iteration of the interior-point algorithm becomes in the case of the two-stage SAA SP problem:

$$\begin{bmatrix} Q_0 + D_0^2 & & & & T_0^T & \frac{1}{N} T_1^T & \dots & \frac{1}{N} T_N^T \\ & \frac{1}{N} (Q_1 + D_1^2) & & & 0 & \frac{1}{N} W_1^T & \dots & 0 \\ & & \ddots & & \vdots & \vdots & \ddots & \vdots \\ & & & \frac{1}{N} (Q_N + D_N^2) & 0 & 0 & \dots & \frac{1}{N} W_N^T \\ T_0 & 0 & \dots & 0 & & & & \\ \frac{1}{N} T_1 & \frac{1}{N} W_1 & \dots & 0 & & & & \\ \vdots & \vdots & \ddots & \vdots & & & & \\ \frac{1}{N} T_N & 0 & \dots & \frac{1}{N} W_N & & & & \end{bmatrix} \begin{bmatrix} \Delta x_0 \\ \Delta x_1 \\ \vdots \\ \Delta x_N \\ -\Delta y_0 \\ -\Delta y_1 \\ -\Delta y_2 \\ \vdots \\ -\Delta y_N \end{bmatrix} = r,$$

where $D_i = X_i^{-\frac{1}{2}} U_i^{\frac{1}{2}}$ and r is computed accordingly. From now on we will use the notation $\bar{Q}_i := Q_i + D_i^2$, $i = 0, 1, 2, \dots, N$.

A symmetric permutation of the above system given by

$$\pi = \begin{pmatrix} 1 & 2 & 3 & \dots & N & N+1 & N+2 & N+3 & \dots & 2N+1 & 2N+2 \\ 2N+1 & 1 & 3 & \dots & 2N-3 & 2N-1 & 2N+2 & 2 & \dots & 2N-2 & 2N \end{pmatrix}$$

is performed to obtain the following arrow-shaped linear system that is specific to stochastic programming and suitable for parallelization by using the Schur complement:

$$\begin{bmatrix} \frac{1}{N}\bar{Q}_1 & \frac{1}{N}W_1^T & & & 0 & 0 \\ \frac{1}{N}W_1 & 0 & & & \frac{1}{N}T_1 & 0 \\ & & \ddots & & \vdots & \vdots \\ & & & \frac{1}{N}\bar{Q}_N & \frac{1}{N}W_N^T & 0 & 0 \\ & & & \frac{1}{N}W_N & 0 & \frac{1}{N}T_N & 0 \\ 0 & \frac{1}{N}T_1^T & \dots & 0 & \frac{1}{N}T_N^T & \bar{Q}_0 & T_0^T \\ 0 & 0 & \dots & 0 & 0 & T_0 & 0 \end{bmatrix} \begin{bmatrix} \Delta x_1 \\ -\Delta y_1 \\ \vdots \\ \Delta x_N \\ -\Delta y_N \\ \Delta x_0 \\ -\Delta y_0 \end{bmatrix} = \pi r. \quad (11)$$

Solving this linear system is the main computational effort at each iteration of the interior-point algorithm. When the Schur complement of the lower right 2-by-2 block in the system's matrix is used, a scenario-based decomposition is obtained. The same mechanism is used in [18–20, 46] to achieve the parallelization of the problem. We briefly show how the factorization as well as the forward and back substitutions phases can be parallelized; we refer the reader to the abovementioned papers for more details.

We further simplify the notation and denote by $K_1, K_2, \dots, K_N, K_0$ the diagonal blocks and by B_1, B_2, \dots, B_N the bordering blocks of the system matrix of (11). Hence

$$K_i := \begin{bmatrix} \frac{1}{N}\bar{Q}_i & \frac{1}{N}W_i^T \\ \frac{1}{N}W_i & 0 \end{bmatrix}, \quad K_0 := \begin{bmatrix} \bar{Q}_0 & W_0^T \\ W_0 & 0 \end{bmatrix},$$

$$B_i := \begin{bmatrix} 0 & 0 \\ \frac{1}{N}T_i & 0 \end{bmatrix}, \quad i = 1, 2, \dots, N.$$

Let $\Delta z_i := [\Delta x_i \ -\Delta y_i]^T$, $i = 0, 1, \dots, N$, $\Delta z := [\Delta z_1 \ \dots \ \Delta z_N \ \Delta z_0]^T$, and let r be of the form $[r_1 \ \dots \ r_N \ r_0]^T$.

With the new notation, the matrix of the linear system (11) becomes

$$K := \begin{bmatrix} K_1 & & & B_1 \\ & \ddots & & \vdots \\ & & K_N & B_N \\ B_1^T & \dots & B_N^T & K_0 \end{bmatrix}. \quad (12)$$

Since K is symmetric, it can be factorized as LDL^T , where L is a unit lower triangular matrix and D is a diagonal matrix. One can easily verify that L and D have the following particular structures,

$$L = \begin{bmatrix} L_1 & & & \\ \vdots & \ddots & & \\ & & L_N & \\ L_{N1} & \dots & L_{NN} & L_c \end{bmatrix}, \quad D = \begin{bmatrix} D_1 & & & \\ & \ddots & & \\ & & D_N & \\ & & & D_c \end{bmatrix},$$

where

$$L_i D_i L_i^T = K_i, \quad i = 1, \dots, N, \quad (13)$$

$$L_{Ni} = B_i^T L_i^{-T} D_i^{-1}, \quad i = 1, \dots, N, \quad (14)$$

$$L_c D_c L_c^T = C := K_0 - \sum_{i=1}^N B_i^T K_i^{-1} B_i. \quad (15)$$

Observe that C defined above is the Schur complement of the first-stage Hessian block K_0 in the entire Hessian matrix K .

Consequently, the following steps are needed to solve $K\Delta z = r$:

$$w_i = L_i^{-1}r_i, \quad i = 1, \dots, N, \quad (16)$$

$$w_0 = L_c^{-1} \left(r_0 - \sum_{i=1}^N L_{Ni}w_i \right), \quad (17)$$

$$v_i = D_i^{-1}w_i, \quad i = 0, 1, \dots, N, \quad (18)$$

$$\Delta z_0 = L_c^{-1}v_0, \quad (19)$$

$$\Delta z_i = L_i^{-T}(v_i - L_{Ni}\Delta z_0), \quad i = 1, \dots, N. \quad (20)$$

These operations can be divided into four phases:

- *Factorization*: (13) and (15);
- *Forward substitution*: steps (16) and (17);
- *Diagonal solve*: step (18);
- *Back substitution*: steps (19) and (20).

The matrices L_{Ni} from (13) are not explicitly computed and stored, and the products from (17) and (20) are computed ad hoc.

3 The preconditioning

In this section we first introduce and analyze the spectral properties of a preconditioner for the symmetric positive definite $(1, 1)$ block of the Schur complement matrix C . The preconditioner approximates the $(1, 1)$ block of C by using a subset of the existing scenarios; it is called “stochastic preconditioner”. We then present two approaches for iteratively solving $C\Delta z_0 = r$. The first approach algebraically incorporates the stochastic preconditioner into a symmetric indefinite preconditioner for C . The resulting preconditioner is known in the literature as the constraint preconditioner. The second approach uses the stochastic preconditioner with a projected conjugate gradient method to solve a reduced linear system.

The expression (15) of the Schur complement matrix C can be rewritten as

$$C = \begin{bmatrix} \bar{Q}_0 & T_0^T \\ T_0 & 0 \end{bmatrix} - \frac{1}{N} \sum_{i=1}^N \begin{bmatrix} 0 & T_i^T \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \bar{Q}_i & W_i^T \\ W_i & 0 \end{bmatrix}^{-1} \begin{bmatrix} 0 & 0 \\ T_i & 0 \end{bmatrix}. \quad (21)$$

We use the following well-known matrix inversion identity to express the Schur complement matrix C in terms of the data of the original problem.

Lemma 1 *For any nonsingular matrix A and any full-rank matrix B the following matrix identity holds:*

$$\begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix}^{-1} = \begin{bmatrix} A^{-1} - A^{-1}B^T(BA^{-1}B^T)^{-1}BA^{-1} & A^{-1}B^T(BA^{-1}B^T)^{-1} \\ (BA^{-1}B^T)^{-1}BA^{-1} & -(BA^{-1}B^T)^{-1} \end{bmatrix}.$$

Proof By direct verification.

By applying the above identity to each of the inner blocks of the summation from (21), C can be written as

$$C = \begin{bmatrix} \bar{Q}_0 + \frac{1}{N} \sum_{i=1}^N T_i^T \left(W_i \bar{Q}_i^{-1} W_i^T \right)^{-1} T_i & T_0^T \\ T_0 & 0 \end{bmatrix}. \quad (22)$$

3.1 The stochastic preconditioner

Let us denote the (1, 1) block from (22) by S_N and observe that

$$S_N = \left(Q_0 + D_0^2 \right) + \frac{1}{N} \sum_{i=1}^N M_i,$$

where

$$M_i = T_i^T \left(W_i \left(Q_i + D_i^2 \right)^{-1} W_i^T \right)^{-1} T_i.$$

Let $\xi_{k_1}, \xi_{k_2}, \dots, \xi_{k_n}$ be a uniformly distributed subsequence from $\xi_1, \xi_2, \dots, \xi_N$ with $n \ll N$, and denote

$$S_n = \left(Q_0 + D_0^2 \right) + \frac{1}{n} \sum_{i=1}^n M_{k_i}.$$

We propose S_n as a preconditioner for S_N .

In the remainder of this section we show that the probability that the eigenvalues of the preconditioned matrix $S_n^{-1} S_N$ are outside an ϵ -ball centered at 1 exponentially approaches 0 as n increases. In other words, the eigenvalues of the preconditioned matrix cluster around 1 “exponentially” with the number of subsamples. The fact that the eigenvalues cluster around 1 is obvious since for $n = N$ the preconditioned matrix is exactly the identity matrix. The exponential rate of clustering will make the preconditioner effective since n is taken to be much more smaller than N for the sake of parallel efficiency.

In particular we regard the entries of matrices M_i , $i = 1, \dots, N$ as independent identically distributed (IID) realizations of the corresponding entries of the random matrix

$$M(\omega) := T^T(\omega) \left(W(\omega) \left(Q(\omega) + D^2(\omega) \right)^{-1} W^T(\omega) \right)^{-1} T(\omega), \quad (23)$$

and use the theory of large deviations to obtain the exponential convergence bounds. One can see that $M(\omega)$ depends on the barrier parameter $\mu > 0$ through the diagonal matrix $D(\omega)$. The analysis of this section is pointwise with μ ; that is, the convergence results hold for any $\mu > 0$, but the exponential bounds change with μ .

We now present the assumptions on the random data $\xi(\omega)$ that are needed to ensure that the entries of $M(\omega)$ are bounded random variables. We first assume the boundedness of the random data of the two-stage SP problem:

(A1) There exists $U > 0$ such that $\|\xi(\omega)\| < U$, for any $\omega \in \Omega$.

Second, we require the singular values of the recourse matrix $W(\omega)$ to be bounded away from 0 and the singular values of the technology matrix $T(\omega)$ to stay bounded uniformly with ω :

(A2) There exists $\sigma_m^W > 0$ such that

$$0 < \sigma_m^W \leq \sigma_{\min}(W(\omega)), \text{ for any } \omega \in \Omega.$$

(A3) There exists $\sigma_M^T \in \mathbb{R}$ such that

$$\sigma_{\max}(T(\omega)) \leq \sigma_M^T \text{ for any } \omega \in \Omega.$$

Another, obvious in some way, assumption is that the quadratic objective matrix $Q(\omega)$ should also stay bounded uniformly with ω , namely

(A4) There exists $\lambda_M^Q > 0$ such that

$$\lambda_{\max}(Q(\omega)) \leq \lambda_M^Q \text{ for any } \omega \in \Omega.$$

Note that in the case of a discrete or compact probability space Ω , Assumptions (A2), (A3), and (A4) follow from Assumption (A1). Therefore, the conditions under which we analyze the preconditioner hold for a large class of applications.

Furthermore, the homeomorphism of the mapping defining the central path (see [23] or [35]) and Assumption (A4) allow us, for a fixed barrier parameter $\mu > 0$, to consider $x = x(\omega)$ and $u = u(\omega)$ as bounded random vectors. Consequently, we can regard the diagonal matrices D_i , $i = 1, \dots, n$ as realizations of the random matrix $D(\omega) = X^{-\frac{1}{2}}(\omega)U^{\frac{1}{2}}(\omega)$. The boundedness of $D(\omega)$ follows from the fact that $x(\omega)$ must be positive since it satisfies $xu = \mu e$ with $\mu > 0$.

As we mentioned in Section 2, path-following algorithms follow the central path approximately, and the $x(\omega)$ and $u(\omega)$ are actually situated in a neighborhood of the central path. The proof of the randomness and boundedness of $D(\omega)$ for such $x(\omega)$ and $u(\omega)$ is intractable in the case of Mehrotra's algorithm because this method does not use an explicit neighborhood of the central path. However, the algorithm maintains the positiveness of the x and u by backing off by a fixed factor whenever the boundary of the positive orthant is reached. Therefore, we assume that the entries of $D(\omega)$ are bounded random variables:

(A5) For a fixed $\mu > 0$, there exists positive numbers d_m^μ and d_M^μ such that

$$0 < d_m^\mu \leq \left[D^2(\omega) \right]_{ii} \leq d_M^\mu, \text{ for any } i \text{ and } \omega \in \Omega.$$

We now present the spectral results that are used in the proof of Lemma 3.

Lemma 2 *Let $Q \in \mathbb{R}^{n \times n}$ be a symmetric positive definite matrix and $A \in \mathbb{R}^{m \times n}$ be a full-rank matrix. The following affirmations hold:*

(i) *The extreme eigenvalues of Q are characterized by*

$$\lambda_{\min}(Q) = \min_{x \neq 0} \frac{x^T Q x}{\|x\|^2} \text{ and } \lambda_{\max}(Q) = \max_{x \neq 0} \frac{x^T Q x}{\|x\|^2} = \|Q\|_2.$$

(ii) *The largest singular value of A is characterized by*

$$\|A\|_2 = \max_{x \neq 0} \frac{\|Ax\|}{\|x\|} = \sigma_{\max}(A) = \max_{y \neq 0} \frac{\|A^T y\|}{\|y\|} = \|A^T\|_2.$$

When $m \leq n$, the smallest singular value of A satisfies

$$\sigma_{\min}(A) = \min_{x \neq 0} \frac{\|A^T x\|}{\|x\|}.$$

(iii) In the case $m \leq n$, the extreme eigenvalues of AQA^T satisfy

$$\lambda_{\min}(Q) \sigma_{\min}^2(A) \leq \lambda_{\min}(AQA^T) \leq \lambda_{\max}(AQA^T) \leq \lambda_{\max}(Q) \sigma_{\max}^2(A).$$

Proof (i) A simple proof can be obtained based on the symmetric Schur decomposition of Q (page 393 in [15]) or based on the Courant-Fischer minimax theorem (also in [15], page 394).

(ii) See [15], page 71.

(iii) For any $y \in \mathbb{R}^m$, from (i) and (ii) we can write

$$\begin{aligned} y^T AQA^T y &\geq \lambda_{\min}(Q) \|A^T y\|^2 \geq \lambda_{\min}(Q) \sigma_{\min}^2(A) \|y\|^2 \text{ and} \\ y^T AQA^T y &\leq \lambda_{\max}(Q) \|A^T y\|^2 \leq \lambda_{\max}(Q) \sigma_{\max}^2(A) \|y\|^2. \end{aligned}$$

These inequalities and (i) applied to AQA^T imply that $\lambda_{\min}(Q) \sigma_{\min}^2(A)$ is a lower bound for $\lambda_{\min}(AQA^T)$ and $\lambda_{\max}(Q) \sigma_{\max}^2(A)$ is an upper bound for $\lambda_{\max}(AQA^T)$, and hence inequalities from (iii) hold.

Lemma 3 Under Assumptions (A1)-(A5) the matrix

$$M(\omega) = T^T(\omega) \left(W(\omega) \left(Q(\omega) + D^2(\omega) \right)^{-1} W^T(\omega) \right)^{-1} T(\omega)$$

is well defined, and each of its entries is a bounded random variable.

Proof Since the entries of $D^2(\omega)$ are positive and the $Q(\omega)$ is positive semidefinite, we have that (see [15, Theorem 8.1.5])

$$d_m^\mu \leq \lambda_{\min} \left(Q(\omega) + D^2(\omega) \right) \leq \lambda_{\max} \left(Q(\omega) + D^2(\omega) \right) \leq d_M^\mu + \lambda_M^Q, \forall \omega, \quad (24)$$

which implies that the matrix $Q(\omega) + D^2(\omega)$ is invertible. From (A2) we obtain that $W(\omega) \left(Q(\omega) + D^2(\omega) \right)^{-1} W^T(\omega)$ is also invertible. This shows that $M(\omega)$ is well defined and its components are random variables.

We now prove the componentwise boundedness. Observe that (24) also allows us to deduce that

$$\lambda_{\min} \left(\left(Q(\omega) + D^2(\omega) \right)^{-1} \right) \geq \frac{1}{d_M^\mu + \lambda_M^Q}, \forall \omega. \quad (25)$$

Then by (iii) of Lemma 2 the following bound exists for the smallest eigenvalue of the matrix $W(\omega) \left(Q(\omega) + D^2(\omega) \right)^{-1} W^T(\omega)$,

$$\lambda_{\min} \left(W(\omega) \left(Q(\omega) + D^2(\omega) \right)^{-1} W^T(\omega) \right) \geq \frac{(\sigma_m^W)^2}{d_M^\mu + \lambda_M^Q}, \forall \omega.$$

Therefore,

$$\lambda_{\max} \left(\left(W(\omega) \left(Q(\omega) + D^2(\omega) \right)^{-1} W^T(\omega) \right)^{-1} \right) \leq \frac{d_M^\mu + \lambda_M^Q}{(\sigma_m^W)^2}, \forall \omega. \quad (26)$$

By the submultiplicative property of the 2-norms and based on (i) and (ii) of Lemma 2 and the bound (26), the 2-norm of $M(\omega)$ can be bounded as follows

$$\begin{aligned} \|M(\omega)\|_2 &\leq \|T^T(\omega)\|_2 \left\| \left(W(\omega) (Q(\omega) + D^2(\omega))^{-1} W^T(\omega) \right) \right\|_2 \|T(\omega)\|_2 \\ &\leq \frac{(\sigma_M^T)^2}{(\sigma_m^W)^2} (d_M^\mu + \lambda_M^Q), \forall \omega. \end{aligned}$$

The boundedness of the entries of $M(\omega)$ follows immediately by writing

$$|[M(\omega)]_{ij}| = |e_i^T M(\omega) e_j| \leq \|M(\omega)\| \leq \frac{(\sigma_M^T)^2}{(\sigma_m^W)^2} (d_M^\mu + \lambda_M^Q), \forall \omega,$$

and hence the lemma is proved.

Lemma 3 allows us to regard S_N as the sample average of $S := \mathbb{E}[S]$, where S is the random matrix $S(\omega) := (Q_0 + D_0^2) + M(\omega)$. By the law of large numbers $S_N \rightarrow S$ w.p. 1 as $N \rightarrow \infty$. Similarly, S_n can be seen as a sample average of S , although one obtained from fewer samples than S_N .

We now show that the eigenvalues of $S_N^{-1}S$ cluster exponentially fast with N around 1. The result will also hold for the matrix $S_n^{-1}S_N$ since S_n is a sample average of S_N with respect to the empirical distribution given by ξ_1, \dots, ξ_N .

The boundedness of $S(\omega)$ also follows from Lemma 3. We simplify the notation by letting $L > 0$ (depending on $\lambda_m^Q, \sigma_m^W, \sigma_M^T$, and d_M^μ as per Lemma 3) such that

$$|[S(\omega)]_{ij}| < L, \text{ for any } \omega \in \Omega.$$

Then, by the Hoeffding's inequality (see Chapter 7 of [42]), we have that for any $\epsilon > 0$,

$$Pr(|[S_k]_{ij} - S_{ij}| \geq \epsilon) \leq 2 \exp\left(-\frac{\epsilon^2 N}{2L^2}\right). \quad (27)$$

The following lemma characterizes the distance between the components of the identity matrix and the preconditioned matrix $S_N^{-1}S$. We denote by p the size of the matrices S and S_N and by $\|S^{-1}\|_{max}$ the largest magnitude of the components of S .

Lemma 4 For any $\epsilon > 0$,

$$Pr\left(\left|[I - S^{-1}S_N]_{ij}\right| \geq \epsilon\right) \leq 2p \exp\left(-\frac{\epsilon^2 N}{2p^2 L^2 \|S^{-1}\|_{max}^2}\right). \quad (28)$$

Proof Since $I - S^{-1}S_N = S^{-1}(S - S_N)$, we have

$$\left|[I - S^{-1}S_N]_{ij}\right| = \left|\sum_{k=1}^p [S^{-1}]_{ik} [S - S_N]_{kj}\right| \leq \|S^{-1}\|_{max} \sum_{k=1}^p |[S - S_N]_{kj}|,$$

which allows us to write

$$Pr\left(\left|[I - S^{-1}S_n]_{ij}\right| \geq \epsilon\right) \leq Pr\left(\sum_{k=1}^p |[S - S_N]_{kj}| \geq \frac{\epsilon}{\|S^{-1}\|_{max}}\right). \quad (29)$$

In bounding the right term in the above inequality we use the following inequality holding for any random variables Y_1, Y_2, \dots, Y_p :

$$\Pr\left(\sum_{k=1}^p Y_k \geq a\right) \leq \sum_{k=1}^p \Pr(Y_k \geq \frac{a}{p}). \quad (30)$$

For a proof see [42, Inequality (7.101)]. Then (29) can be further bounded as follows:

$$\begin{aligned} \Pr\left(\left|[I - S^{-1}S_n]_{ij}\right| \geq \epsilon\right) &\leq \sum_{k=1}^p \Pr\left(\left|[S - S_N]_{kj}\right| \geq \frac{\epsilon}{p\|S^{-1}\|_{max}}\right) \\ &\leq \sum_{k=1}^p 2 \exp\left(-\frac{\epsilon^2 N}{2p^2 L^2 \|S^{-1}\|_{max}^2}\right) \text{ (by (27))} \\ &= 2p \exp\left(-\frac{\epsilon^2 N}{2p^2 L^2 \|S^{-1}\|_{max}^2}\right), \end{aligned}$$

which proves the thesis.

Lemma 5 *If the symmetric matrices $A, B \in \mathbb{R}^{p \times p}$ satisfy*

$$\Pr(|A_{ij} - B_{ij}| \geq \epsilon) \leq c \exp(-C\epsilon^2), \forall i, j \in \{1, 2, \dots, p\}$$

for some constants $c > 0$ and $C > 0$, then the eigenvalues of A and B are characterized by

$$\Pr(|\lambda_k(A) - \lambda_k(B)| \geq \epsilon) \leq cp^2 \exp\left(\frac{-C\epsilon^2}{p^2}\right).$$

Proof By the Wielandt-Hoffman theorem [15] we have that

$$\sum_k (\lambda_k(A) - \lambda_k(B))^2 \leq \|A - B\|_F^2. \quad (31)$$

Furthermore, we can write

$$\begin{aligned} \Pr(|\lambda_k(A) - \lambda_k(B)| \geq \epsilon) &= \Pr\left(|\lambda_k(A) - \lambda_k(B)|^2 \geq \epsilon^2\right) \\ &\leq \Pr\left(\sum_{k=1}^p (\lambda_k(A) - \lambda_k(B))^2 \geq \epsilon^2\right) \\ &\leq \Pr\left(\|A - B\|_F^2 \geq \epsilon^2\right) \text{ (by (31))} \\ &= \Pr\left(\sum_{i,j=1}^p |A_{ij} - B_{ij}|^2 \geq \epsilon^2\right) \\ &\leq \sum_{i,j=1}^p \Pr\left(|A_{ij} - B_{ij}|^2 \geq \epsilon^2/p^2\right) \text{ (by (30))} \\ &= \sum_{i,j=1}^p \Pr\left(|A_{ij} - B_{ij}| \geq \epsilon/p\right) \leq cp^2 \exp\left(\frac{-C\epsilon^2}{p^2}\right), \end{aligned}$$

which completes the proof.

The following result shows that the eigenvalues of the preconditioned matrix $S_N^{-1}S$ cluster around 1 exponentially with the sample size N .

Theorem 1 For any $\epsilon > 0$ the eigenvalues of $S_N^{-1}S$ satisfy

$$Pr \left(|\lambda(S_N^{-1}S) - 1| \geq \epsilon \right) \leq 2p^3 \exp \left(-\frac{N(1 - \frac{1}{1+\epsilon})^2}{2p^4 L^2 \|S^{-1}\|_{max}^2} \right).$$

Proof By applying Lemma 5 to the matrices I and $S^{-1}S_N$ and using the inequality (28) given in Lemma 4, we can write

$$Pr \left(|\lambda(S^{-1}S_N) - 1| \geq \epsilon \right) \leq 2p^3 \exp \left(-\frac{\epsilon^2 N}{2p^4 L^2 \|S^{-1}\|_{max}^2} \right).$$

Since $\lambda(S^{-1}S_N) = 1/\lambda(S_N^{-1}S)$, this inequality can be transformed to

$$Pr \left(\frac{1}{1+\epsilon} \leq \lambda(S_N^{-1}S) \leq \frac{1}{1-\epsilon} \right) \leq 2p^3 \exp \left(-\frac{\epsilon^2 N}{2p^4 L^2 \|S^{-1}\|_{max}^2} \right)$$

or

$$Pr \left(1 - \frac{\epsilon}{1+\epsilon} \leq \lambda(S_N^{-1}S) \leq 1 + \frac{\epsilon}{1-\epsilon} \right) \leq 2p^3 \exp \left(-\frac{\epsilon^2 N}{2p^4 L^2 \|S^{-1}\|_{max}^2} \right).$$

Since $\epsilon/(1+\epsilon) < \epsilon/(1-\epsilon)$, we obtain

$$Pr \left(1 - \frac{\epsilon}{1+\epsilon} \leq \lambda(S_N^{-1}S) \leq 1 + \frac{\epsilon}{1+\epsilon} \right) \leq 2p^3 \exp \left(-\frac{\epsilon^2 N}{2p^4 L^2 \|S^{-1}\|_{max}^2} \right).$$

By taking $\epsilon \leftarrow \epsilon/(1+\epsilon) = 1 - 1/(1+\epsilon)$ we obtain

$$Pr \left(|\lambda(S_N^{-1}S) - 1| \geq \epsilon \right) \leq 2p^3 \exp \left(-\frac{N(1 - \frac{1}{1+\epsilon})^2}{2p^4 L^2 \|S^{-1}\|_{max}^2} \right),$$

which completes the proof of the theorem.

As we mentioned before, the bound of Theorem 1 holds for the preconditioner of interest, $S_n^{-1}S_N$, since S_n is the average approximation of S_N in the empirical distribution given by ξ_1, \dots, ξ_N .

3.2 The constraint preconditioner

A popular preconditioning technique for saddle-point matrices such as C is the class of constraint preconditioners [25], that incorporates an existing preconditioner of the $(1, 1)$ block together with the $(1, 2)$ and $(2, 1)$ blocks of the original saddle-point system.

The constraint preconditioner corresponding to C is

$$M = \begin{bmatrix} S_n & T_0^T \\ T_0 & 0 \end{bmatrix}. \quad (32)$$

The spectral analysis from [2] (see proof of Theorem 4.1 of the paper) reveals that the preconditioned matrix $M^{-1}C$ has an eigenvalue at 1 with order of multiplicity $2m_0$ and $n_0 - m_0$ real eigenvalues satisfying

$$0 < \lambda_{min}(S_n^{-1}S_N) \leq \lambda(M^{-1}C) \leq \lambda_{max}(S_n^{-1}S_N),$$

where n_0 and m_0 are the number of rows of S_n and T_0 , respectively. Hence the constraint preconditioner (32) possesses the same exponential clustering of the eigenvalues around 1 as the stochastic preconditioner we introduced in the previous section.

Algorithm DSC for Process $p \in \mathcal{P}$

Factorization phase

1. For each $i \in \mathcal{N}_p$ factorize $L_i D_i L_i^T = K_i$;
2. Compute $C_p = - \sum_{i \in \mathcal{N}_p} B_i^T K_i^{-1} B_i$, reduce $C = \sum_{r \in \mathcal{P}} C_r$
and compute $C = C + K_0$;
- 3*. Factorize the Schur complement $L_C D_C L_C^T = C$;

Back substitution phase

4. For each $i \in \mathcal{N}_p$ solve $w_i = L_i^{-1} r_i$;
5. Compute $\sum_{i \in \mathcal{N}_p} L_{N_i} w_i$ and reduce $\sum_{i=1}^N L_{N_i} w_i$;
- 6*. Solve $w_0 = L_C^{-1} (r_0 - \sum_{i=1}^N L_{N_i} w_i)$;

Diagonal solve phase

7. For each $i \in \mathcal{N}_p$ solve $v_i = D_i^{-1} w_i$;
- 8*. Solve $v_0 = D_C^{-1} w_0$;

Forward substitution phase

- 9*. Solve $\Delta z_0 = L_C^{-T} v_0$;
 10. For each $i \in \mathcal{N}_p$ solve $\Delta z_i = L_i^{-T} (v_i - L_{N_i} \Delta z_0)$.
-

4 Parallel implementation

The linear solve $K \Delta x = r$ from Section 2 is suitable to a scenario-based parallelization. While the factorizations and triangular solves involving the second-stage variable are fully parallelizable, the factorization and triangular solves involving the first stage variable are not because the Schur complement C given by (15) requires data from all scenarios.

Our parallel implementation distributes both the data and the work corresponding to second-stage variables across multiple processes and uses the Message Passing Interface (MPI) as the underlying mechanism for communication between processes. The Schur complement C is stored and computed on all processes. This approach causes a bottleneck that it is addressed by the preconditioning technique presented in Section 3.

Scenarios are evenly assigned to available computational units by solving a number partitioning problem. More exactly, if $\mathcal{P} = \{1, 2, \dots, P\}$ denotes the set of available computational units, and if there are N scenarios, each with positive load l_i , $i = \{1, 2, \dots, N\}$, the number partitioning problem consists of finding a partition $\mathcal{N}_1, \mathcal{N}_2, \dots, \mathcal{N}_P$ of $\{1, 2, \dots, N\}$ such that the differences between

$$\sum_{i \in \mathcal{N}_1} l_i, \sum_{i \in \mathcal{N}_2} l_i, \dots, \sum_{i \in \mathcal{N}_P} l_i$$

are as small as possible, ideally zero. The load l_i is the wall-clock time needed to perform work associated with scenario i at the previous interior-point iteration. For the first interior-point iteration the load is computed based on the size and the fill-in of the scenario's data.

We first list the algorithm that uses the direct factorization of the Schur complement matrix to solve the linear system $K \Delta z = r$ given in Section 2. We call this algorithm *DSC* (direct Schur complement); it solves a two-stage problem involving N scenarios by using P processes. Once the partition $\mathcal{N}_1, \mathcal{N}_2, \dots, \mathcal{N}_P$ of $\{1, 2, \dots, N\}$ is obtained, any

Algorithm PSC- p ($p = 1, 2, \dots, P$)

Factorization phase

- 1.1. For each $i \in \mathcal{K}_p$ factorize $L_i D_i L_i^T = K_i$;
- 2.1. Compute $C_{\mathcal{K}_p} = - \sum_{i \in \mathcal{K}_p} B_i^T K_i^{-1} B_i$, reduce $C_{\mathcal{K}} = \sum_{r=1}^{P+1} C_{\mathcal{K}_r}$ on process $P+1$ and on process 1;
- 1.2. For each $i \in \mathcal{N}_p$ factorize $L_i D_i L_i^T = K_i$;
- 2.2. Compute $C_{\mathcal{N}_p} = - \sum_{i \in \mathcal{N}_p} B_i^T K_i^{-1} B_i$, reduce $C_{\mathcal{N}} = \sum_{r \in \mathcal{P}} C_r$ and on process $p = 1$ only compute $C = C_{\mathcal{K}} + C_{\mathcal{N}} + K_0$;

Back substitution phase

4. For each $i \in \mathcal{K}_p \cup \mathcal{N}_p$ solve $w_i = L_i^{-1} r_i$;
5. Compute $\sum_{i \in \mathcal{K}_p \cup \mathcal{N}_p} L_{N_i} w_i$ and reduce $\sum_{i=1}^N L_{N_i} w_i$ on process 1;

Iterative preconditioned solve (6*, 8*, 9*)

Krylov solve for $\Delta z_0 = C^{-1}(r_0 - \sum_{i=1}^N L_{N_i} w_i)$ (only process 1);
 Process 1 sends and processes $2, \dots, P$ receive Δz_0 ;

Diagonal solve phase

7. For each $i \in \mathcal{K}_p \cup \mathcal{N}_p$ solve $v_i = D_i^{-1} w_i$;

Forward substitution phase

10. For each $i \in \mathcal{K}_p \cup \mathcal{N}_p$ solve $\Delta z_i = L_i^{-T}(v_i - L_{N_i} \Delta z_0)$.
-

process $p \in \mathcal{P}$ performs the set of operations listed in Algorithm *DSC*. The first-stage operations that are performed on all processes are marked with a “*”.

4.1 Parallel implementation of the preconditioning

The factorization from step 3* and the solves from steps 6*, 8*, and 9* of Algorithm *DSC* represent bottlenecks in the parallel execution flow because all processes must perform them. In particular, the bottleneck caused by the factorization 3* of C adversely affects the parallel scaling of Algorithm *DSC* for the majority of real-life applications. This adverse behavior is caused by the fact that the Schur complement C is dense, and its factorization starts to dominate to overall execution time when the ratio of scenarios per process become small. We address this bottleneck by using iterative methods that make use of the preconditioner S_n introduced and analyzed in Section 3. The factorization of C is completely removed from the parallel execution flow, and the only bottleneck remains in the much less expensive solve phase.

In this new approach, a Krylov-type iterative solver is applied to the system $C \Delta z_0 = r_0 - \sum_{i=1}^N L_{N_i} w_i$, since the factors of C are not available anymore. In general, preconditioned Krylov subspace methods require the inverse of the preconditioner to be applied to a vector at each Krylov iteration. In our case, this step is done cheaply by means of triangular solves with factors of the preconditioner. The factors of the preconditioner are computed by a separate process $P+1$ *in the same time* the other P processes compute the terms of full Schur complement matrix C .

To deal with a heterogeneous set of processes, we implemented a slightly more sophisticated scheduling mechanism than that of Algorithm *DSC*. Namely, let $\mathcal{K} = \{k_1, k_2, \dots, k_n\}$ be the scenarios based on which the preconditioner S_n is constructed,

Algorithm $PSC-(P+1)$ **factorization phase**

- 1.1. For each $i \in \mathcal{K}_{P+1}$ factorize $L_i D_i L_i^T = K_i$;
- 2.1. Compute $C_{\mathcal{K}_{P+1}} = - \sum_{i \in \mathcal{K}_{P+1}} B_i^T K_i^{-1} B_i$, reduce $C_{\mathcal{K}} = \sum_{r=1}^{P+1} C_{\mathcal{K}_r}$ on process $P+1$ and on process 1, and compute $M = C_{\mathcal{K}} + K_0$;
- 3*. Factorize $L_M D_M L_M^T = M$;

Back substitution phase

4. For each $i \in \mathcal{K}_{P+1}$ solve $w_i = L_i^{-1} r_i$;
5. Compute $\sum_{i \in \mathcal{K}_{P+1}} L_{N_i} w_i$ and reduce $\sum_{i=1}^N L_{N_i} w_i$ on process 1;

Iterative preconditioned solve (6*, 8*, 9*)

Apply preconditioner M^{-1} as many times process 1 needs it;
Receive Δz_0 from process 1;

Diagonal solve phase

7. For each $i \in \mathcal{K}_{P+1}$ solve $v_i = D_i^{-1} w_i$;

Forward substitution phase

10. For each $i \in \mathcal{K}_{P+1}$ solve $\Delta z_i = L_i^{-T} (v_i - L_{N_i} \Delta z_0)$.

and let $\bar{\mathcal{N}} := \{1, 2, \dots, N\} \setminus \mathcal{K}$ be the set of remaining scenarios. The scenarios for \mathcal{K} are assigned to all $P+1$ processes, and the scenarios from $\bar{\mathcal{N}}$ are assigned only to processes 1, 2 through P . Once the scenarios from \mathcal{K} are finished and S_n is computed, processes 1, 2 through P continue to compute the terms from the full Schur complement C , while process $P+1$ starts the factorization of the preconditioning. An even distribution of the scenarios is obtained by setting up and solving two distinct number partitioning problems for \mathcal{K} and $\bar{\mathcal{N}}$ similar to the partitioning problem from Algorithm DSC . Let us denote the partitions found by $\mathcal{K}_1, \mathcal{K}_2, \dots, \mathcal{K}_{P+1}$ for \mathcal{K} and by $\mathcal{N}_1, \mathcal{N}_2, \dots, \mathcal{N}_P$ for $\bar{\mathcal{N}}$.

The complete set of operations and the interprocess communication patterns needed to solve the linear system $K \Delta z_0 = r$ are called generically Algorithm PSC (preconditioned Schur complement) and are presented separately in Algorithm $PSC-p$ for processes 1, \dots , P and in Algorithm $PSC-(P+1)$ for process $P+1$.

A comparison of Algorithm $PSC-p$ and Algorithm $PSC-(P+1)$ shows that once the scenarios used in S_n are computed (steps 1.1 and 1.2), processes 1, 2 through P continue to compute the terms from the full Schur complement C (steps 1.2 and 2.1), while process $P+1$ starts the factorization of the preconditioner (step 3*). The largest number n of scenarios included in S_n is chosen such that process $P+1$ completes step 3* before the other P processes finish steps 1.2 and 2.2. Therefore, the bottleneck 3* from Algorithm DSC is removed from the parallel execution flow of Algorithm PSC .

On the other hand, the iterative preconditioned Krylov solve of Algorithm PSC is slightly more expensive than the corresponding solve with the factors of the Schur complement (steps 6*, 8*, and 9*) of Algorithm DSC . More exactly, each Krylov iteration requires the action of the inverse of the preconditioner, which has the same cost as solving with the factors of the Schur complement, namely, $O((n_0 + m_0)^2)$, and a total overhead of $O(l(n_0 + m_0)^2)$ is present in the iterative solve phase of Algorithm PSC , where l is the number of Krylov iterations. However, the savings from the removal of the direct factorization of the Schur matrix are $O((n_0 + m_0)^3)$, therefore, Algorithm PSC is more scalable than Algorithm DSC whenever $l \ll n_0 + m_0$. In the numerical

experiments presented in Section 5, the number l of iterations is no more than 100, usually less than 10, while $n_0 + m_0$ is more than 10,000.

We have implemented a mechanism that decides at runtime whether to use PSC or DSC. This mechanism requires no user intervention, estimates the execution times needed by each of the two methods at the next IPM iteration and uses the one with shorter execution time. The estimates are obtained by timing the linear algebra of the Schur complement matrix (or the preconditioner) and the second-stage subproblems at the IPM iteration that just ended. Such estimates are very reliable in the context of IPMs, staying constant for the most of the iterates, and some of them steadily increasing as the iterates approach the solution (because of the ill-conditioning of IPMs). The implementation uses DSC for the first iteration, unless instructed otherwise, and assumes that PSC needs $l = 10$ Krylov iterations when $\mu > 0.1$ and $l = 10 \text{ floor}(\log_{10}(1/\mu))$ otherwise. Our tests showed that this decision mechanism is robust.

4.2 Implementation of the Krylov subspace iterative methods

The iterative methods we have implemented are BiCGStab [22] and the preconditioned projected conjugate gradient (PPCG) from [21].

We apply BiCGStab to the saddle-point linear system involving the Schur complement matrix C and use the constraint preconditioner M given by (32). The products involving C are computed on process 1, and the products involving M^{-1} are performed remotely on process $P + 1$. At each iteration two of such products are needed.

The PPCG method projects the saddle-point linear system $C\Delta z_0$,

$$\begin{bmatrix} S_N & T_0^T \\ T_0 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} r_1 \\ r_2 \end{bmatrix},$$

onto a basis Z_0 , which spans the null space of T_0 . More exactly, the algorithm starts at a point that satisfies $T_0x = r_2$ and generates directions that lie in the null space of T_0 by using the (preconditioned) projector operator

$$P = Z_0(Z_0^T S_n^{-1} Z_0) Z_0^T.$$

Consequently all the iterates x remain in the affine subspace given by $T_0x = r_2$. Once the iterate x converges, y is found from $(T_0 T_0^T)y = r_1 - T_0x$. The cost of solving for y is small since the factorization is done only once, at the first IPM iteration. An important observation here is that Z_0 does not have to be explicitly computed. Instead, the projection $v = Pu$ can be performed efficiently by solving the linear system

$$\begin{bmatrix} S_n & T_0^T \\ T_0 & 0 \end{bmatrix} \begin{bmatrix} v \\ w \end{bmatrix} = \begin{bmatrix} u \\ 0 \end{bmatrix},$$

that is, by solving with the factors of the *constraint preconditioner* M , not with those of the stochastic preconditioner S_n . The residual update strategy [21] enforced to reduce the effects of roundoff errors occurring in the computation of the projected residual requires an extra solve with the factors of M at each iteration of the PPCG. Therefore, two applications of the inverse of the preconditioner is needed at each PPCG iteration, which makes PPCG and BiCGStab to have about the same cost per one iteration.

5 Computational results

In this section we present the computational results of the DSC and PSC approaches presented in the previous section. We have used *Fusion* a 320-node computing cluster at Argonne National Laboratory using an InfiniBand QDR interconnect with 2 μ sec latency. Each node is dual-socket, quad-core and therefore has 8 cores. Each core operates at 2.53 GHz. A minimum of 16 GB memory is available for each node.

Both the DSC and PSC algorithms were implemented in a parallel interior-point solver for stochastic programming (PIPS) we developed recently. PIPS uses the object-oriented design of OOQP [14] and reuses many of OOQP's classes. The largest problem we have solved with PIPS consists of 28.9 millions variables. The problem is similar to the unit commitment problem described in Section 5.1 but has simplified dynamics and has been replicated artificially for testing purposes. The speedup obtained by PIPS using the DSC approach is shown in Figure 1. Parallel efficiencies of 0.95, 0.92, and 0.77 are reached when the number of cores is increased from 80 to 400, 600, and 1000, respectively. In the case of a building energy system control problem [45], almost perfect scaling was obtained from 10 cores to 50 cores. The ratio of the number of first-stage variables and the number of second-stage variables is small for both problems, and therefore the bottleneck created by the Schur complement does not affect significantly the scalability. However, this is not the case for the problems described in the following section, and the PSC method is able to overcome this limiting behavior as will be shown in Section 5.3.




Fig. 1 Speedup of PIPS using DSC method in solving a large unit commitment problem.

5.1 The test problems

The computational framework presented in [8] integrates the Weather Research and Forecasting (WRF) model in stochastic unit commitment and energy dispatch formulations that account for the wind power in the generation of the electricity. The unit commitment (UC) problem of the above study is a mixed integer linear programming problem and involves a network of 10 thermal generators and 12 wind farms (each of 50 turbines). We solved a relaxation of the problem and obtained the sample average approximation problem by using a set of 30 scenarios for the wind power levels. We denote this instance as *UC1* and use it in the next subsection to study the spectrum of the preconditioner.

The second instance, denoted by *UC2*, was obtained by replicating three times the network of generators of *UC1* and using a set of 120 (resampled) scenarios. The problem has a total of 210,000 variables and 443,000 constraints. The first-stage problem is relatively large when compared to the second-stage problems, with 10,800 variables

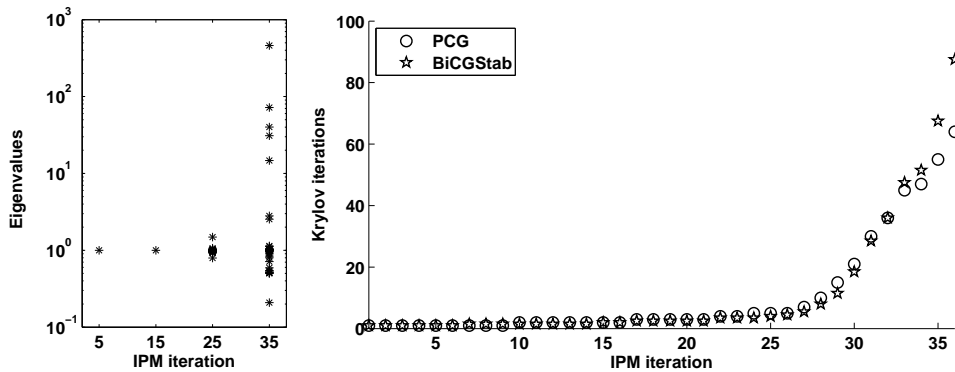


Fig. 2 In the left figure, the 25 largest and smallest eigenvalues of the preconditioned matrix at IPM iterations 5, 15, 25, and 35 are depicted. The right figure shows the number of inner iterations needed to solve the predictor system at each IPM iteration.

and 24,000 constraints and 1,656 variables and 3,494 constraints, respectively. The *UC2* problem is used in subsection 5.3 to investigate and compare the scalability of the DSC and PSC methods.

The above two instances are both linear programming problems ($Q_0 = 0$ and $Q(\omega) = 0$ in (2)). We note that the technology matrix T and the recourse matrix W are fixed (deterministic). The randomness occurs only in the right-hand side $b(\omega)$.

5.2 Quality of preconditioners

The first numerical experiment investigates the quality of the S_n preconditioner when the *UC1* instance is solved. The number of scenarios included in S_n was $n = 5$ out of a total of 30 scenarios. We first look at the eigenvalues of the preconditioned matrix $S_n^{-1}S_N$ at different moments in the solving process: IPM iterations 5, 15, 25, and 35. The smallest and largest 25 eigenvalues of $S_n^{-1}S_N$ are displayed in the left plot of Figure 2. This figure shows a solid clustering of the eigenvalues at 1. The dispersion of the extreme several eigenvalues as the optimization approaches the solution is caused by the adverse effect of the ill-conditioning of the diagonals D_i , $i = 0, \dots, n$ on the average approximation S_n of S_N . Consequently, both BiCGStab and PPCG take just less than 3 iterations for more than half of the outer iterations but need an increasing number of iterations after that. At the end of the IPM iterations, we recorded a maximum of 88 iterations for BiCGStab and 64 iterations for PPCG, which should be seen as very good numbers since the linear system solved with a relative error of 10^{-8} was 2608×2608 and extremely ill-conditioned.

Another important observation is that the number of iterations needed by BiCGStab and PPCG is almost the same for the first 30 IPM iterations. After that PPCG takes advantage of the positive-definiteness of the $(1, 1)$ block and needs fewer iterations than BiCGStab to solve the system within the same accuracy.

5.3 Practical performance of the preconditioners

We use the *UC2* problem to study the scalability of DSC and PSC. The mechanism described in Section 4.1 that decides whether to use DSC or PSC was disabled for the tests presented here. *Strong scaling* is investigated, that is, the same problem was solved with an increasing number of cores. A linear scaling occurs when the execution time decreases linearly as the number of cores increases. *UC2* problem has a very large Schur complement matrix with adversely affects the strong scaling of the DSC, as can be seen in Figure 3. PSC does not exhibit the same limitation and scales linearly for a wide range of cores, 2-20. When using more than 20 cores, however, the factorization of the preconditioner does not finish before the scenarios are done and thus creates its own bottleneck in the parallel execution flow. Consequently, the scaling is no longer linear, as shown in Figure 3.

We have used DSC with 2 cores as a reference for the strong scaling of not only DSC but also PSC. Therefore, Figure 3 also shows which method is faster. For example, the reduction in the solve time obtained by PSC over DSC is 20%, 28%, 33%, and 24% for 10, 15, 20, and 30 cores, respectively.

6 Conclusions

In this paper, we presented a preconditioning technique for Schur complement linear systems arising in the interior-point based parallelization of the two-stage stochastic programming problems with recourse. The spectral analysis based on the Hoeffding's inequality from the theory of large deviations showed that an exponentially fast clustering of the eigenvalues occurs with the number of scenarios incorporated in the preconditioner. The good quality of the preconditioner allows the IPM saddle-point systems to be solved with the same accuracy as the direct solvers with a small number of Krylov iterations.

Numerical experiments involving large instances of a stochastic unit commitment problem indicate good practical performance. The small number of Krylov iterations needed to solve the (preconditioned) Schur complement linear systems significantly reduces the bottleneck caused by the direct factorization of the Schur complement matrix. Consequently, the preconditioning technique exhibits a linear strong scaling for a reasonably wide range of processes. This range can be further extended by preconditioning with more than one process. While the algorithms we presented in this paper will not change in this case, the implementation must parallelize the factorization of the preconditioner. We plan to follow this approach and to implement it in PIPS.

The preconditioning algorithm PSC can also be used when solving an SAA T -stage stochastic programming, $T > 2$. In this case, the IPM saddle-point linear systems have an arrow shape *nested* structure; that is, each of the diagonal blocks K_i , $i = 1, \dots, N$, of K from (12) have the arrow shape of K , recursively defined for $T > 3$. The Schur

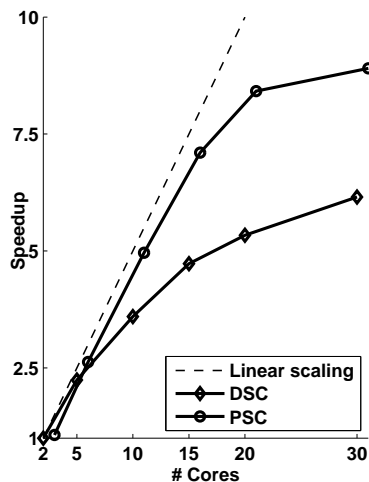


Fig. 3 Strong scaling of DSC and PSC using P and $P + 1$ cores, respectively ($P = \{2, 5, 10, 15, 20, 30\}$).

complement technique we described in this paper can be then used for each of the first $T - 1$ stages to decompose the problem [17]. The scalability of the DSC method may be affected by stages with a disproportionately large number of variables, and the PSC method should be used for such stages. We also note that a mixed DSC-PSC approach for multistage problems is robust, in the sense that the implementation can decide at runtime for each stage whether it is profitable or not to employ the preconditioning based on the decision mechanism presented in Section 4.1.

We point out that the performance of the preconditioner improves in terms of the number of Krylov iterations if a regularization technique of the saddle-point system is used [1]. Moreover, the preconditioner will perform better on convex quadratic programming problems than on linear programming problem because of the presence of the positive (semi)definite quadratic term that ameliorates (some of) the ill-conditioning of the Schur complement.

Acknowledgments

The authors are grateful to John Birge for discussions on the issues in the manuscript. This work was supported by the U.S. Department of Energy through contract DE-AC02-06CH11357.

References

1. Altman, A., Gondzio, J.: Regularized symmetric indefinite systems in interior-point methods for linear and quadratic optimization. *Optimization Methods and Software* **11**(1-4), 275–302 (1999)
2. Bergamaschi, L., Gondzio, J., Zilli, G.: Preconditioning indefinite systems in interior point methods for optimization. *Computational Optimization and Applications* **28**(2), 149–171 (2004)
3. Birge, J.R.: Current trends in stochastic programming computation and applications. Tech. rep., Department of Industrial and Operations Engineering, University of Michigan, Ann Harbour, Michigan (1995)
4. Birge, J.R., Chen, X., Qi, L.: A stochastic Newton method for stochastic quadratic programs with recourse. Tech. rep., Applied Mathematics Preprint AM94/9, School of Mathematics, the University of New South Wales (1995)
5. Birge, J.R., Holmes, D.F.: Efficient solution of two stage stochastic linear programs using interior point methods. *Computational Optimization and Applications* **1**, 245–276 (1992)
6. Birge, J.R., Louveaux, F.: *Introduction to stochastic programming*. Springer-Verlag, New York (1997)
7. Birge, J.R., Qi, L.: Computing block-angular Karmarkar projections with applications to stochastic programming. *Manage. Sci.* **34**(12), 1472–1479 (1988)
8. Constantinescu, E.M., Zavala, V.M., Rocklin, M., Lee, S., Anitescu, M.: A computational framework for uncertainty quantification and stochastic optimization in unit commitment with wind power generation. *IEEE Transactions on Power Systems*, to appear (2010)
9. Czyzyk, J., Mehrotra, S., Wright, S.J.: PCx user guide. Tech. Rep. OTC 96/01, Optimization Technology Center, Argonne National Laboratory and Northwestern University (1996)
10. Dantzig, G.B., Infanger, G.: Large-scale stochastic linear programs – Importance sampling and Benders decomposition. In: *Computational and Applied Mathematics, I*, pp. 111–120. North-Holland, Amsterdam (1992)
11. Dollar, H.S., Gould, N.I.M., Schilders, W.H.A., Wathen, A.J.: Implicit-factorization preconditioning and iterative solvers for regularized saddle-point systems. *SIAM Journal on Matrix Analysis and Applications* **28**(1), 170–189 (2006)
12. Dollar, H.S., Wathen, A.J.: Approximate factorization constraint preconditioners for saddle-point matrices. *SIAM Journal on Scientific Computing* **27**(5), 1555–1572 (2006)

13. Ermoliev, Y.M.: Stochastic quasigradient methods. In: Numerical techniques for stochastic optimization, *Springer Ser. Comput. Math.*, vol. 10, pp. 141–185. Springer, Berlin (1988)
14. Gertz, E.M., Wright, S.J.: Object-oriented software for quadratic programming. *ACM Transactions on Mathematical Software* **29**(1), 58–81 (2003)
15. Golub, G.H., Van Loan, C.F.: *Matrix Computations*, 3rd edn. The Johns Hopkins University Press (1996)
16. Gondzio, J.: HOPDM (version 2.12) - a fast LP solver based on a primal-dual interior point method. *European Journal of Operational Research* **85**, 221–225 (1995)
17. Gondzio, J., Grothey, A.: Direct solution of linear systems of size 10^9 arising in optimization with interior point methods. In: PPAM, pp. 513–525 (2005)
18. Gondzio, J., Grothey, A.: Parallel interior-point solver for structured quadratic programs: Application to financial planning problems. *Annals of Operations Research* **152**(1), 319–339 (2007)
19. Gondzio, J., Grothey, A.: Exploiting structure in parallel implementation of interior point methods for optimization. *Computational Management Science* **6**(2), 135–160 (2009)
20. Gondzio, J., Sarkissian, R.: Parallel interior point solver for structured linear programs. *Mathematical Programming* **96**, 561–584 (2000)
21. Gould, N.I.M., Hribar, M.E., Nocedal, J.: On the solution of equality constrained quadratic programming problems arising in optimization. *SIAM Journal on Scientific Computing* **23**(4), 1376–1395 (2001)
22. Greenbaum, A.: *Iterative Methods for Solving Linear Systems*. SIAM Series on Frontiers in Applied Mathematics, Philadelphia, PA (1997)
23. Güler, O.: Existence of interior points and interior paths in nonlinear monotone complementarity problems. *Math. Oper. Res.* **18**(1), 128–147 (1993)
24. Higle, J.L., Sen, S.: Stochastic decomposition: An algorithm for two-stage linear programs with recourse. *Mathematics of Operations Research* **16**, 650–669 (1991)
25. Keller, C., Gould, N.I.M., Wathen, A.J.: Constraint preconditioning for indefinite linear systems. *SIAM J. Matrix Anal. Appl.* **21**(4), 1300–1317 (2000)
26. King, A.J.: An implementation of the Lagrangian finite-generation method. In: Numerical techniques for stochastic optimization, *Springer Ser. Comput. Math.*, vol. 10, pp. 295–311. Springer, Berlin (1988)
27. Linderoth, J., Wright, S.J.: Decomposition algorithms for stochastic programming on a computational grid. *Comput. Optim. Appl.* **24**(2-3), 207–250 (2003)
28. Lukšan, L., Vlček, J.: Indefinitely preconditioned inexact Newton method for large sparse equality constrained nonlinear programming problems. *Numerical Linear Algebra with Applications* **5**(3), 219–247 (1998)
29. Lustig, I.J., Marsten, R.E., Shanno, D.F.: On implementing Mehrotra’s predictor–corrector interior-point method for linear programming. *SIAM Journal on Optimization* **2**(3), 435–449 (1992)
30. McKay, M.D., Beckman, R.J., Conover, W.J.: A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* **21**, 239–245 (1979)
31. Mehrotra, S.: On the implementation of a primal-dual interior point method. *SIAM Journal on Optimization* **2**(4), 575–601 (1992)
32. Mehrotra, S., Ozevin, M.G.: Decomposition-based interior point methods for two-stage stochastic semidefinite programming. *SIAM J. on Optimization* **18**(1), 206–222 (2007)
33. Mehrotra, S., Ozevin, M.G.: Decomposition based interior point methods for two-stage stochastic convex quadratic programs with recourse. *Oper. Res.* **57**(4), 964–974 (2009)
34. Mehrotra, S., Ozevin, M.G.: On the implementation of interior point decomposition algorithms for two-stage stochastic conic programs. *SIAM Journal on Optimization* **19**(4), 1846–1880 (2009)
35. Monteiro, R.D.C., Pang, J.S.: Properties of an interior-point mapping for mixed complementarity problems. *Mathematics of Operations Research* **21**(3), 629–654 (1996)
36. Murphy, M.F., Golub, G.H., Wathen, A.J.: A note on preconditioning for indefinite linear systems. *SIAM Journal on Scientific Computing* **21**(6), 1969–1972 (2000)
37. Pflug, G.C., Halada, L.: A note on the recursive and parallel structure of the birge and qi factorization for tree structured linear programs. *Comput. Optim. Appl.* **24**(2-3), 251–265 (2003). DOI <http://dx.doi.org/10.1023/A:1021810125060>
38. Rockafellar, R.T., Wets, R.J.B.: A Lagrangian finite generation technique for solving linear-quadratic problems in stochastic programming. *Math. Programming Stud.* **28**, 63–93 (1986). Stochastic programming 84. II

39. Rockafellar, R.T., Wets, R.J.B.: Scenarios and policy aggregation in optimization under uncertainty. *Mathematics of Operations Research* **16**(1), 119–147 (1991)
40. Rosa, C.H., Ruszczyński, A.: On augmented Lagrangian decomposition methods for multi-stage stochastic programs. *Ann. Oper. Res.* **64**, 289–309 (1996). *Stochastic programming, algorithms and models* (Lillehammer, 1994)
41. Ruszczyński, A.: A regularized decomposition method for minimizing a sum of polyhedral functions. *Mathematical Programming* **35**, 309–333 (1986)
42. Shapiro, A., Dentcheva, D., Ruszczyński, A.: *Lectures on Stochastic Programming: Modeling and Theory*. MPS/SIAM Series on Optimization 9, Philadelphia, PA (2009)
43. Van Slyke, R., Wets, R.J.: L-shaped linear programs with applications to control and stochastic programming. *SIAM Journal on Applied Mathematics* **17**, 638–663 (1969)
44. Wright, S.J.: *Primal-Dual Interior-Point Methods*. Society for Industrial and Applied Mathematics, Philadelphia, PA (1997)
45. Zavala, V.M., Constantinescu, E.M., Krause, T., Anitescu, M.: On-line economic optimization of energy systems using weather forecast information. *Journal of Process Control* **19**, 1725–1736 (2009). DOI 10.1016/j.jprocont.2009.07.004
46. Zavala, V.M., Laird, C.D., Biegler, L.T.: Interior-point decomposition approaches for parallel solution of large-scale nonlinear parameter estimation problems. *Chemical Engineering Science* **63**(19), 4834–4845 (2008)
47. Zhang, D., Zhang, Y.: A Mehrotra-type predictor-corrector algorithm with polynomiality and Q -subquadratic convergence. *Ann. Oper. Res.* **62**, 131–150 (1996)
48. Zhang, Y.: Solving large-scale linear programs by interior-point methods under the Matlab environment. Tech. Rep. TR96-01, University of Maryland Baltimore County (1996)
49. Zhao, G.: A log-barrier method with Benders decomposition for solving two-stage stochastic linear programs. *Mathematical Programming* **90**(3), 507–536 (2001)

(To be removed before publication) The submitted manuscript has been created by the University of Chicago as Operator of Argonne National Laboratory (“Argonne”) under Contract No. DE-AC02-06CH11357 with the U.S. Department of Energy. The U.S. Government retains for itself, and others acting on its behalf, a paid-up, nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.