European Journal of Operational Research 231 (2013) 263–273

Contents lists available at SciVerse ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Continuous Optimization

Improving an interior-point approach for large block-angular problems by hybrid preconditioners

Silvana Bocanegra [a,b], Jordi Castro [b,*], Aurelio R.L. Oliveira [c]

[a] Dept. de Estatística e Informática, Universidade Federal Rural de Pernambuco, R. Dom Manoel de Medeiros, s/n, Dois Irmãos, 52171-900 Recife, Brazil
[b] Dept. of Statistics and Operations Research, Universitat Politècnica de Catalunya, Campus Nord, Office C5 203, Jordi Girona 1-3, 08034 Barcelona, Catalonia, Spain
[c] Dept. of Applied Mathematics, University of Campinas, R. Sérgio Buarque de Holanda, 651, 13083-859 Campinas, SP, Brazil

A R T I C L E I N F O

Article history:
Received 31 July 2012
Accepted 7 April 2013
Available online 15 April 2013

Keywords:
Interior-point methods
Large-scale optimization
Preconditioned conjugate gradient
Structured problems

A B S T R A C T

The computational time required by interior-point methods is often dominated by the solution of linear systems of equations. An efficient specialized interior-point algorithm for primal block-angular problems has been used to solve these systems by combining Cholesky factorizations for the block constraints and a conjugate gradient based on a power series preconditioner for the linking constraints. In some problems this power series preconditioner resulted to be inefficient on the last interior-point iterations, when the systems became ill-conditioned. In this work this approach is combined with a splitting preconditioner based on LU factorization, which works well for the last interior-point iterations. Computational results are provided for three classes of problems: multicommodity flows (oriented and nonoriented), minimum-distance controlled tabular adjustment for statistical data protection, and the minimum congestion problem. The results show that, in most cases, the hybrid preconditioner improves the performance and robustness of the interior-point solver. In particular, for some block-angular problems the solution time is reduced by a factor of 10.

1. Introduction

Many important large-scale optimization problems exhibit a block-angular structure. Applications are found in fields such as control and planning, network flows, stochastic linear programming, and statistical data protection. Several interior-point methods have been devised to solve these structured problems [5,7,12,16,25]. These specialized algorithms exploit the particular structure of the constraints matrix, and some were implemented for parallel environments [5,25]. The efficiency of interior-point methods critically depends of the linear system solver used at each iteration to compute the Newton direction. Such systems are often written in a symmetric indefinite form, known as the *augmented system*. They can also be reduced to a smaller positive definite form, the *normal equations*. Techniques based on direct and iterative solvers can be applied for their solution. For some classes of large scale problems the use of direct methods becomes prohibitive due to storage and time limitations, whereas iterative linear solvers with appropriate preconditioners may be more efficient.

The efficient interior-point algorithm for primal block-angular problems of [15] solved the normal equations in two stages: Cholesky factorizations for the block constraints and a Preconditioned Conjugate Gradient (PCG) for the linking constraints. The purpose of PCG is to avoid solving the system associated to the complicating linking constraints by Cholesky factorizations, in an attempt to make the problem block separable. The preconditioner is obtained by truncating an infinite power series that approximates the inverse of the system to be solved. For some difficult primal block-angular problems this approach outperformed state-of-the-art commercial solvers [16]. However, in some problems, systems become very ill-conditioned as the optimal solution is reached, and then PCG provides slow and inaccurate solutions. It was shown [16] that the efficiency of this approach depends on the spectral radius—in $[0,1)$—of a certain matrix which appears in the definition of the preconditioner (which is itself related to the Schur complement of the normal equations). Spectral radius close to 1 degrades the performance of the preconditioner. When PCG gives inaccurate solutions, the code implemented in [15] switches to the solution of the normal equations by a Cholesky factorization, which may be prohibitive for large-scale problems.

In order to yield a reliable and efficient interior-point method based just on iterative solvers we introduce a hybrid and adaptive scheme for solving the normal equations. On the first interior-point iterations the normal equations are solved using the Cholesky-PCG approach of [15] outlined above. When the system associated to linking constraints becomes ill-conditioned, the normal equations are solved by a PCG using the splitting preconditioner of [29,30], instead of switching to a direct solver. The splitting preconditioner

* Corresponding author. Tel.: +34 934015854.
E-mail addresses: silvana@deinfo.ufrpe.br (S. Bocanegra), jordi.castro@upc.edu (J. Castro), aurelio@ime.unicamp.br (A.R.L. Oliveira).

0377-2217/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2013.04.007

is a generalization of the tree preconditioner of [33] for large-scale minimum cost network flow problems. Based on a LU factorization, the splitting preconditioner was specially tailored for the last interior-point iterations, when the systems are ill-conditioned. We developed a new and efficient criterion to identify when (i.e., at which interior-point iteration) to switch between iterative solvers. This criterion is based on both the Ritz values of the matrix that appears in the definition of the power series preconditioner, and the number of PCG iterations needed at each interior-point iteration. The Ritz values are approximations of the eigenvalues of a matrix; they will be used to estimate the spectral radius, which measures the efficiency of the power series preconditioner. An implementation of this new approach, combining the power series and the splitting preconditioners, was applied to three classes of primal block-angular instances [15]: multicommodity flows (oriented and nonoriented), minimum-distance controlled tabular adjustment for statistical data protection, and the minimum congestion problem. As it will be shown, the hybrid approach was more efficient than the power series preconditioner in many block-angular problems. Other hybrid approaches combining interior-point and combinatorial algorithms have been used for some type of networks flows problems [21].

This paper is organized as follows. In Section 2 we recall the basic ideas of interior-point methods for primal block-angular problems using the power series preconditioner. The new hybrid approach is described in Section 3, together with an outline of the splitting preconditioner, and a description of the switching criterion between preconditioners. Numerical experiments are shown in Section 4. The effect of different regularization parameters for the splitting preconditioner are also discussed in Section 4. Finally, in Section 5 the conclusions are drawn and further developments are suggested.

2. The interior-point algorithm for primal block-angular problems

One of the most efficient interior-point methods for some classes of block-angular problems was initially developed for multicommodity flows [12] and later extended for general primal block-angular problems [15]. This method considers the following general formulation of a block-angular problem:

$$\begin{array}{ll} \min & \sum_{i=o}^{k} ({c^i}^T x^i + {x^i}^T Q_i x^i) \\ \text{s.t.} & \begin{bmatrix} N_1 & & & & \\ & N_2 & & & \\ & & \ddots & & \\ & & & N_k & \\ L_1 & L_2 & \dots & L_k & I \end{bmatrix} \begin{bmatrix} x^1 \\ x^2 \\ \vdots \\ x^k \\ x^0 \end{bmatrix} = \begin{bmatrix} b^1 \\ b^2 \\ \vdots \\ b^k \\ b^0 \end{bmatrix} \\ & 0 \leqslant x^i \leqslant u^i \quad i = 1, \dots, k. \end{array} \tag{1}$$

Matrices $N_i \in \mathbb{R}^{m_i \times n_i}$ and $L_i \in \mathbb{R}^{l \times n_i},\ i = 1, \dots, k$, define, respectively, the block and linking constraints, k being the number of blocks. Vectors $x^i \in \mathbb{R}^{n_i},\ i = 1, \dots, k$, are the variables for each block. $x^0 \in \mathbb{R}^l$ are the slacks of the linking constraints. $b^i \in \mathbb{R}^{m_i},\ i = 1, \dots, k$ is the right-hand-side vector for each block of constraints, whereas $b^0 \in \mathbb{R}^l$ is for the linking constraints. The upper bounds for each group of variables are defined by u^i, $i = 1, \dots, k$. This formulation considers the general form of linking constraints $b^0 - u^0 \leqslant \sum_{i=1}^k L_i x^i \leqslant b^0$. $c^i \in \mathbb{R}^{n_i}$ and $Q_i \in \mathbb{R}^{n_i \times n_i}$, $i = 1, \dots, k$, are the linear and quadratic costs for each group of variables. We also consider linear and quadratic costs $c^0 \in \mathbb{R}^l$ and $Q_0 \in \mathbb{R}^{l \times l}$ for the slacks. We restrict our considerations to the separable case where Q_i, $i = 0, \dots, k$, are diagonal positive semidefinite matrices.

Problem (1) can be written in standard form as

$$\begin{array}{ll} \min & c^T x + \frac{1}{2} x^T Q x \\ \text{s.t.} & Ax = b \\ & 0 \leqslant x \leqslant u \end{array} \tag{2}$$

where $A \in \mathbb{R}^{m \times n}$ $(m = l + \sum_{i=1}^k m_i,\quad n = l + \sum_{i=1}^k n_i$ and $m \leqslant n)$, $Q \in \mathbb{R}^{n \times n}, b \in \mathbb{R}^m$ and $c, x, u \in \mathbb{R}^n$. Replacing inequalities in (2) by a logarithmic barrier with parameter $\mu > 0$, we obtain the logarithmic barrier problem

$$\begin{array}{ll} \min & B(x, \mu) \triangleq c^T x + \frac{1}{2} x^T Q x - \mu \sum_{i=1}^n \ln(x_i) - \mu \sum_{i=1}^n \ln(u_i - x_i) \\ \text{s.t.} & Ax = b. \end{array} \tag{3}$$

The first order KKT optimality conditions for the logarithmic barrier problem—or equivalently, the perturbed KKT-μ conditions for (2)—are

$$\begin{aligned} & Ax = b, \\ & A^T y - Qx + z - w = c, \\ & XZe = \mu e, \\ & (U - X)We = \mu e, \\ & (z, w) > 0, \quad u > x > 0, \end{aligned} \tag{4}$$

where $y \in \mathbb{R}^m, z \in \mathbb{R}^n, w \in \mathbb{R}^n$ are, respectively, the Lagrange multipliers of constraints $Ax = b$, $x \geqslant 0$ and $x \leqslant u$, $X, Z, U, W \in \mathbb{R}^{n \times n}$ are diagonal matrices made up of vectors x, z, u, w, and $e \in \mathbb{R}^n$ is a vector of 1's. The first two sets of equations of (4) impose, respectively, primal and dual feasibility, while the remaining two impose perturbed complementarity. The set of primal–dual solutions $\mathcal{C} = \{(x_\mu, y_\mu, z_\mu, w_\mu), \mu > 0\}$ of (4) is known as the *central path*. Primal–dual path-following interior-point algorithms approximately follow the central path by applying Newton's method to the nonlinear system of Eq. (4), reducing the barrier parameter μ at each iteration. When $\mu \to 0$ these solutions converge to the optimal solution of the original problem. Full details can be found in [37]. The Newton direction is obtained by solving a linear system in variables Δx, Δy, Δz and Δw. In practice, variables Δz and Δw are eliminated and the system reduces to the indefinite augmented system form

$$\begin{bmatrix} -\Theta^{-1} & A^T \\ A & \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} r \\ r_b \end{bmatrix}, \tag{5}$$

where Θ and r are defined as

$$\Theta = (Q + S^{-1}W + X^{-1}Z)^{-1} \quad r = r_c + S^{-1} r_{sw} - X^{-1} r_{xz}, \tag{6}$$

and $S = U - X$. Eliminating Δx from the first group of equations system (5) is reduced to the smaller positive definite normal equations

$$(A \Theta A^T) \Delta y = r_b + A \Theta r = g. \tag{7}$$

For separable quadratic optimization problems Q and Θ are diagonal, and normal equations are usually the preferred choice for computing the Newton direction.

2.1. Normal equations for block-diagonal problems

The performance of interior-point methods relies on the efficient solution of either (5) or (7). For block-angular problems (1) matrices A and Θ have a special structure. The interior-point algorithm used in this work [12,15] solves the normal Eq. (7) by exploiting the block decomposition of $A \Theta A^T$: