

SECOND-ORDER CONE PROGRAMMING METHODS FOR TOTAL VARIATION-BASED IMAGE RESTORATION*

DONALD GOLDFARB[†] AND WOTAO YIN[†]

Abstract. In this paper we present optimization algorithms for image restoration based on the total variation (TV) minimization framework of Rudin, Osher, and Fatemi (ROF). Our approach formulates TV minimization as a second-order cone program which is then solved by interior-point algorithms that are efficient both in practice (using nested dissection and domain decomposition) and in theory (i.e., they obtain solutions in polynomial time). In addition to the original ROF minimization model, we show how to apply our approach to other TV models, including ones that are not solvable by PDE-based methods. Numerical results on a varied set of images are presented to illustrate the effectiveness of our approach.

Key words. image denoising, total variation, second-order cone programming, interior-point methods, nested dissection, domain decomposition

AMS subject classifications. 68U10, 65K10, 90C25, 90C51

DOI. 10.1137/040608982

1. Introduction. Digital images, no matter what their source, usually contain noise. Consequently, a fundamental problem in image processing is the restoration or denoising of such images. Early methods for doing this were based on least squares and had the unfortunate property of either smoothing edges or creating spurious oscillations near edges, i.e., the well-known ringing phenomenon. In [19], Rudin, Osher, and Fatemi (ROF) proposed a method based on minimizing the *total variation* (TV) of the image. Specifically, they proposed solving the minimization problem

$$(1.1) \quad \begin{aligned} \min \quad & \int_{\Omega} \|\nabla u\| \, dx \\ \text{subject to (s.t.)} \quad & u + v = f, \\ & \int_{\Omega} |v|^2 \, dx \leq \sigma^2, \end{aligned}$$

where $\|\nabla u\|$ denotes the Euclidean norm of ∇u at each x , Ω is an open (typically rectangular) domain in \mathbb{R}^2 , $f : \Omega \rightarrow \mathbb{R}$ is a given noisy image in $L^2(\Omega)$, u is the computed estimate of the original image, and σ^2 is an estimate of the variance of the noise in the image f . Minimizing the TV of u allows u to have discontinuities; hence edges and important features in the original image are preserved by the ROF approach.

Rather than solving the constrained minimization problem (1.1), ROF and subsequent researchers also considered the unconstrained minimization problem

$$(1.2) \quad \min \int_{\Omega} \|\nabla u\| \, dx + \lambda \int_{\Omega} |f - u|^2 \, dx,$$

*Received by the editors May 24, 2004; accepted for publication (in revised form) February 4, 2005; published electronically October 31, 2005. This research was supported by NSF grant DMS 01-04282, ONR grant N00014-03-1-0514, and DOE grant GE-FG01-92ER-25126.

<http://www.siam.org/journals/sisc/27-2/60898.html>

[†]Department of IEOR, Columbia University, New York, NY 10027 (goldfarb@columbia.edu, wy2002@columbia.edu).

which yields the same solution as (1.1) for a suitable choice of the Lagrange multiplier λ (see, e.g., [6]). To solve (1.2), ROF proposed using an artificial time stepping method to compute the steady state solution u of the parabolic system $\frac{\partial u}{\partial t} = g(u)$ with initial condition $u = f$ and homogeneous Neumann boundary conditions, where

$$(1.3) \quad g(u) \equiv \frac{1}{2\lambda} \nabla \cdot \left(\frac{\nabla u}{\|\nabla u\|} \right) + f - u.$$

Note that $g(u) = 0$ is the Euler–Lagrange equation for (1.2), which is a necessary and sufficient condition for u to be a solution of the convex minimization problem (1.2). Other algorithms that have been proposed for solving problem (1.2) can be found in [5, 7, 9, 12, 13, 17, 23].

We propose here a direct method for solving (1.1) based on the observation that a discretized version of this convex optimization problem can be reformulated as a second-order cone program (SOCP). One can also reformulate a discretized version of the convex unconstrained minimization problem (1.2) as an SOCP. Our approach is to solve the SOCP representation of (1.1) by an interior-point method taking advantage of the specific structure imparted to the SOCP by the image restoration problem and the TV-based denoising method. Of the many methods that have been proposed for solving (1.1) (or (1.2)) or related problems, it is most closely related to the one proposed by Chan, Golub, and Mulet in [9] based on the interior-point algorithm of Andersen et al. [3].

In the next section we provide a formal definition of an SOCP, present the discretized version of problem (1.1) that we will address, and introduce some notation. In section 3 we give an SOCP formulation of the ROF-based image restoration problem. We also present a dual formulation of problem (1.1) that is essentially equivalent to the dual formulation proposed by Chambolle [5] and show how it can be reformulated as an SOCP. In section 4 we discuss implementational aspects that enable interior-point methods to solve these problems efficiently in practice. In particular, we show how nested dissection can be used to reduce storage and computational effort. A detailed analysis is presented in the appendix. In theory the solution is obtained in polynomial time. Section 5 contains our numerical results. It also contains a discussion of how “overlapping domain decomposition” can be used to further reduce storage and run times. Finally, in section 6 we discuss how our approach can be applied to other image processing models. In particular, we show that it can be applied to the model introduced by Meyer in [15] that is not solvable by PDE-based methods.

2. Notation and preliminaries. The images we discuss in this paper are defined as two-dimensional $n \times n$ matrices in \mathbb{R}^{n^2} . We restrict our discussion to square domains only for the sake of simplicity. Let f be an observed image, which, in real applications, is a noisy version of an original image u . By introducing the noise term v , f and u have the following relation:

$$(2.1) \quad f_{i,j} = u_{i,j} + v_{i,j} \quad \text{for } i, j = 1, \dots, n.$$

$f_{i,j}$ and $u_{i,j}$ are the values of the observed image and the original image, respectively, at pixel (i, j) . Let ∂^+ be the discrete differential operator defined by

$$(2.2) \quad \partial^+ u_{i,j} \stackrel{\text{def}}{=} \left((\partial_x^+ u)_{i,j}, (\partial_y^+ u)_{i,j} \right),$$

where

$$(2.3) \quad \begin{aligned} (\partial_x^+ u)_{i,j} &\stackrel{\text{def}}{=} u_{i+1,j} - u_{i,j} && \text{for } i = 1, \dots, n-1, j = 1, \dots, n, \\ (\partial_y^+ u)_{i,j} &\stackrel{\text{def}}{=} u_{i,j+1} - u_{i,j} && \text{for } i = 1, \dots, n, j = 1, \dots, n-1. \end{aligned}$$

In addition, the differentials on the image domain boundary, $(\partial_x^+ u)_{n,j}$ for $j = 1, \dots, n$ and $(\partial_y^+ u)_{i,n}$ for $i = 1, \dots, n$, are defined to be zero. Consequently, they do not contribute to the total variation, which in the discrete case is defined by

$$(2.4) \quad TV(u) \stackrel{\text{def}}{=} \sum_{1 \leq i,j \leq n} \|\partial^+ u_{i,j}\|,$$

where $\|\cdot\|$ denotes the Euclidean norm, i.e., $\|\partial^+ u_{i,j}\| = ((\partial_x^+ u)_{i,j})^2 + ((\partial_y^+ u)_{i,j})^2)^{1/2}$.

Applying our definitions of f , u , v , and $\partial^+ u$ to the ROF restoration model (1.1), we obtain the finite-dimensional minimization problem

$$(2.5) \quad \begin{aligned} \min \quad & \sum_{1 \leq i,j \leq n} \|\partial^+ u_{i,j}\| \\ \text{s.t.} \quad & u + v = f, \\ & \|v\| \leq \sigma. \end{aligned}$$

Problem (2.5) can be reformulated as an SOCP, the *standard form* of which we now define. The vector of variables $\mathbf{x} \in \mathbb{R}^n$ in a standard form SOCP is composed of subvectors $\mathbf{x}_i \in \mathbb{R}^{n_i}$ —i.e., $\mathbf{x} \equiv (\mathbf{x}_1; \mathbf{x}_2; \dots; \mathbf{x}_r)$ —where $n = n_1 + n_2 + \dots + n_r$ and each subvector \mathbf{x}_i must lie in an elementary *second-order cone* of dimension n_i

$$\mathcal{K}^{n_i} \equiv \{\mathbf{x}_i = (x_i^0; \bar{\mathbf{x}}_i) \in \mathbb{R} \times \mathbb{R}^{n_i-1} : \|\bar{\mathbf{x}}_i\| \leq x_i^0\}.$$

DEFINITION 2.1. *The standard form SOCP (see, e.g., [1]) is*

$$(2.6) \quad \begin{aligned} \min \quad & \mathbf{c}_1^\top \mathbf{x}_1 + \dots + \mathbf{c}_r^\top \mathbf{x}_r \\ \text{s.t.} \quad & A_1 \mathbf{x}_1 + \dots + A_r \mathbf{x}_r = \mathbf{b}, \\ & \mathbf{x}_i \in \mathcal{K}^{n_i}, \quad \text{for } i = 1, \dots, r, \end{aligned}$$

where $\mathbf{c}_i \in \mathbb{R}^{n_i}$ and $A_i \in \mathbb{R}^{m \times n_i}$ for $i = 1, \dots, r$ and $\mathbf{b} \in \mathbb{R}^m$.

Since a one-dimensional second-order cone corresponds to a semi-infinite ray, the above standard form SOCP can accommodate nonnegative variables. In fact if all cones \mathcal{K}_i are one-dimensional, then the above SOCP is just a standard form linear program.

If we define \mathcal{K} to be the Cartesian product $\mathcal{K} = \mathcal{K}^{n_1} \times \dots \times \mathcal{K}^{n_r}$ and the vector \mathbf{c} and the matrix A by $\mathbf{c} = (\mathbf{c}_1; \dots; \mathbf{c}_r)$, and $A = (A_1, \dots, A_r)$, then the above SOCP can be written compactly as

$$(2.7) \quad \begin{aligned} \min \quad & \mathbf{c}^\top \mathbf{x} \\ \text{s.t.} \quad & A\mathbf{x} = \mathbf{b}, \\ & \mathbf{x} \in \mathcal{K} \end{aligned}$$

A crucial aspect of the above SOCP is that the cone \mathcal{K} is pointed, closed, convex, and *self-dual*—i.e., the *dual cone* $\mathcal{K}^* \equiv \{\mathbf{z} : \mathbf{x}^\top \mathbf{z} \geq 0 \ \forall \mathbf{x} \in \mathcal{K}\} = \mathcal{K}$ as is the cone corresponding to the nonnegative orthant in \mathbb{R}^n . This similarity between SOCPs and linear programs is more than superficial; as is the case for linear programs, SOCPs can be solved in polynomial time by interior-point methods. This is the approach we propose in this paper.

3. Second-order cone programming formulation.

3.1. Primal formulation. By introducing $t_{i,j}$ and the relations $((\partial_x^+ u)_{i,j})^2 + ((\partial_y^+ u)_{i,j})^2 \leq (t_{i,j})^2$ for each pixel $i, j = 1, \dots, n$, and including equations (2.3) that define $(\partial_x^+ u)_{i,j}$ and $(\partial_y^+ u)_{i,j}$, problem (2.5) can be reformulated as

$$\begin{aligned}
 (3.1) \quad & \min \sum_{1 \leq i,j \leq n} t_{i,j} \\
 & \text{s.t. } u_{i,j} + v_{i,j} = f_{i,j} \quad \text{for } i, j = 1, \dots, n, \\
 & \quad -(\partial_x^+ u)_{i,j} + (u_{i+1,j} - u_{i,j}) = 0 \quad \text{for } i = 1, \dots, n-1, j = 1, \dots, n, \\
 & \quad -(\partial_y^+ u)_{i,j} + (u_{i,j+1} - u_{i,j}) = 0 \quad \text{for } i = 1, \dots, n, j = 1, \dots, n-1, \\
 & \quad (\partial_x^+ u)_{n,j} = 0 \quad \text{for } j = 1, \dots, n, \\
 & \quad (\partial_y^+ u)_{i,n} = 0 \quad \text{for } i = 1, \dots, n, \\
 & \quad v_0 = \sigma, \\
 & \quad (t_{i,j}; (\partial_x^+ u)_{i,j}; (\partial_y^+ u)_{i,j}) \in \mathcal{K}^3 \quad \text{for } i, j = 1, \dots, n, \\
 & \quad (v_0; v) \in \mathcal{K}^{n^2+1},
 \end{aligned}$$

where, for $i, j = 1, \dots, n$, $u_{i,j}$, $(\partial_x^+ u)_{i,j}$, $(\partial_y^+ u)_{i,j}$, $v_{i,j}$, and $t_{i,j}$ are treated as variables and $f_{i,j}$ for $i, j = 1, \dots, n$ and σ as constants.

Although solving for u is our ultimate goal, in the above formulation u can be eliminated and $(\partial_x^+ u)_{n,j} = 0$ for $j = 1, \dots, n$ and $(\partial_y^+ u)_{i,n} = 0$ for $i = 1, \dots, n$ can be removed; hence, (3.1) simplifies to

$$\begin{aligned}
 (3.2) \quad & \min \sum_{1 \leq i,j \leq n, (i,j) \neq (n,n)} t_{i,j} \\
 & \text{s.t. } (\partial_x^+ u)_{i,j} + (v_{i+1,j} - v_{i,j}) = f_{i+1,j} - f_{i,j} \quad \text{for } i = 1, \dots, n-1, j = 1, \dots, n, \\
 & \quad (\partial_y^+ u)_{i,j} + (v_{i,j+1} - v_{i,j}) = f_{i,j+1} - f_{i,j} \quad \text{for } i = 1, \dots, n, j = 1, \dots, n-1, \\
 & \quad v_0 = \sigma, \\
 & \quad (t_{i,j}; (\partial_x^+ u)_{i,j}; (\partial_y^+ u)_{i,j}) \in \mathcal{K}^3 \quad \text{for } i, j = 1, \dots, n-1, \\
 & \quad (t_{n,j}; (\partial_y^+ u)_{n,j}) \in \mathcal{K}^2 \quad \text{for } j = 1, \dots, n-1, \\
 & \quad (t_{i,n}; (\partial_x^+ u)_{i,n}) \in \mathcal{K}^2 \quad \text{for } i = 1, \dots, n-1, \\
 & \quad (v_0; v) \in \mathcal{K}^{n^2+1}.
 \end{aligned}$$

After solving the above problem, u can be recovered via $u_{i,j} = f_{i,j} - v_{i,j}$ for $i, j = 1, \dots, n$.

Problem (3.2) is a standard form SOCP. It has $4n^2 - 2n$ variables, $2n(n-1) + 1$ equality constraints, $(n-1)^2$ three-dimensional and $2(n-1)$ two-dimensional second-order cone constraints, and one large second-order cone constraint.

3.2. Alternative (dual) formulations. Recalling the definition (2.2) of $\partial^+ u_{i,j}$, let us define $\partial^+ u \equiv \{\partial^+ u_{i,j}\} \in \mathbb{R}^{n^2} \times \mathbb{R}^{n^2}$. If we further define $q_{i,j} \equiv (q_{i,j}^1, q_{i,j}^2)$ and

$(\operatorname{div} q) \equiv \{(\operatorname{div} q)_{i,j}\}$, where $(\operatorname{div} q)_{i,j} = (\partial_x^- q^1)_{i,j} + (\partial_y^- q^2)_{i,j}$ and

$$(3.3) \quad \begin{aligned} (\partial_x^- q^1)_{i,j} &\stackrel{\text{def}}{=} q_{i,j}^1 - q_{i-1,j}^1 && \text{for } i = 2, \dots, n-1, j = 1, \dots, n, \\ (\partial_x^- q^1)_{1,j} &\stackrel{\text{def}}{=} q_{1,j}^1 && \text{for } j = 1, \dots, n, \\ (\partial_x^- q^1)_{n,j} &\stackrel{\text{def}}{=} -q_{n-1,j}^1 && \text{for } j = 1, \dots, n, \\ (\partial_y^- q^2)_{i,j} &\stackrel{\text{def}}{=} q_{i,j}^2 - q_{i,j-1}^2 && \text{for } i = 1, \dots, n, j = 2, \dots, n-1, \\ (\partial_y^- q^2)_{i,1} &\stackrel{\text{def}}{=} q_{i,1}^2 && \text{for } i = 1, \dots, n, \\ (\partial_y^- q^2)_{i,n} &\stackrel{\text{def}}{=} -q_{i,n-1}^2 && \text{for } i = 1, \dots, n-1, \end{aligned}$$

then $q \in \mathbb{R}^{n^2} \times \mathbb{R}^{n^2}$ and

$$(3.4) \quad \langle q, \partial^+ u \rangle = \sum_{i,j} (q_{i,j}^1 (\partial_x^+ u)_{i,j} + q_{i,j}^2 (\partial_y^+ u)_{i,j}) = \sum_{i,j} -(\operatorname{div} q)_{i,j} u_{i,j} = \langle -\operatorname{div} q, u \rangle.$$

Let z^* denote the optimal value of

$$(3.5) \quad \begin{aligned} \min \quad & \sum_{1 \leq i,j \leq n} \|\partial^+ u_{i,j}\| \\ \text{s.t.} \quad & \|u - f\|^2 \leq \sigma^2. \end{aligned}$$

Using the fact that $\|a\| = \max_{\|b\| \leq 1} a^\top b$ (an immediate consequence of the Cauchy–Schwarz inequality), if we define the set

$$S \equiv \{q \in \mathbb{R}^{n^2} \times \mathbb{R}^{n^2} : \|q_{i,j}\| \equiv \|(q_{i,j}^1, q_{i,j}^2)\| \leq 1 \ \forall i, j\},$$

then

$$(3.6) \quad z^* = \min_{\|u-f\|^2 \leq \sigma^2} \max_{q \in S} \langle q, \partial^+ u \rangle$$

$$(3.7) \quad = \min_{\|u-f\|^2 \leq \sigma^2} \max_{q \in S} \langle -\operatorname{div} q, u \rangle$$

$$(3.8) \quad = \max_{q \in S} \min_{\|u-f\|^2 \leq \sigma^2} \langle -\operatorname{div} q, u \rangle.$$

From the inner minimization, u must satisfy $u = f + \frac{\operatorname{div} q}{\|\operatorname{div} q\|} \sigma$. Hence, substituting this expression for u and letting $p = -q$ yields the following maximization problem:

$$(3.9) \quad z^* = \max_{p \in S} \{\langle \operatorname{div} p, f \rangle - \sigma \|\operatorname{div} p\|\}.$$

In [5], Chambolle derived a dual formulation of a discrete formulation of problem (1.2), which is the Lagrangian version of (3.5), and proposed a method for solving this formulation using a PDE approach. A related dual formulation is presented in [21] and [4]. We now show how to derive Chambolle's formulation in a direct manner. Specifically, Chambolle considered

$$(3.10) \quad \tilde{z}^* = \min_u \left\{ \frac{\|u - f\|^2}{2\lambda} + \sum_{1 \leq i,j \leq n} \|\partial^+ u_{i,j}\| \right\},$$

which, by arguments identical to those used to derive (3.8), is equivalent to

$$(3.11) \quad \tilde{z}^* = \frac{1}{2\lambda} \max_{q \in S} \left\{ \min_u \{ \|u - f\|^2 + \langle -2\lambda \operatorname{div} q, u \rangle \} \right\}.$$

Since $u = f + \lambda \operatorname{div} q$ solves the inner minimization problem, substituting this expression for u yields

$$(3.12) \quad \tilde{z}^* = \frac{1}{2\lambda} \max_{q \in S} \{ \|\lambda \operatorname{div} q\|^2 + \langle -2\lambda \operatorname{div} q, f + \lambda \operatorname{div} q \rangle \}$$

$$(3.13) \quad = \frac{1}{2\lambda} \max_{q \in S} \{ -\|\lambda \operatorname{div} q + f\|^2 + \|f\|^2 \},$$

which, letting $p = -q$, is equivalent to

$$(3.14) \quad \tilde{z}^* = \frac{1}{2\lambda} \|f\|^2 - \frac{1}{2\lambda} \min_{p \in S} \{ \|\lambda \operatorname{div} p - f\|^2 \}.$$

This minimization problem is identical to problem (8) in [5].

In fact problem (3.9) is equivalent to problem (3.14) for an appropriate choice of σ in terms of λ . For this equivalence to hold, problem (3.5) and its Lagrangian version (3.10) must be equivalent. Hence, for a fixed q , the solutions to the inner minimization problem in (3.8) and (3.11) must be equal; i.e.,

$$(3.15) \quad u = f + \sigma \frac{\operatorname{div} q}{\|\operatorname{div} q\|} = f + \lambda \operatorname{div} q.$$

Since $p = -q$, this gives

$$(3.16) \quad \sigma = \lambda \|\operatorname{div} q\| = \lambda \|\operatorname{div} p\|.$$

Note that adding the term $\|u - f\|^2 / 2\lambda$ to the objective in (3.9) and then using (3.15) and (3.16) to substitute for $u - f$ and σ yields

$$\frac{1}{2\lambda} \max_{p \in S} \{ 2\lambda \langle \operatorname{div} p, f \rangle - \lambda^2 \|\operatorname{div} p\|^2 \},$$

which is easily seen to be equivalent to (3.14).

Next, we present the SOCP formulation of the problem $\min_{q \in S} \{ \|\lambda \operatorname{div} p - f\|^2 \}$. Following the definition of S , the discretization of the above problem gives the following SOCP:

$$(3.17) \quad \begin{aligned} \min \quad & w_0 \\ \text{s.t.} \quad & \lambda(\operatorname{div} q)_{i,j} - w_{i,j} = f_{i,j} \quad \text{for } i, j = 1, \dots, n, \\ & (w_0; (w_{i,j})) \in \mathcal{K}^{n^2+1}, \\ & (q_{i,j}^0; q_{i,j}^1; q_{i,j}^2) \in \mathcal{K}^3 \quad \text{for } i, j = 1, \dots, n-1, \\ & (q_{n,j}^0; q_{n,j}^2) \in \mathcal{K}^2 \quad \text{for } j = 1, \dots, n-1, \\ & (q_{i,n}^0; q_{i,n}^1) \in \mathcal{K}^2 \quad \text{for } i = 1, \dots, n-1, \\ & q_{i,j}^0 = 1 \quad \text{for } i, j = 1, \dots, n, \end{aligned}$$

in which $(\operatorname{div} q)_{i,j}$ is explicitly given by $(\operatorname{div} q)_{i,j} = (\partial_x^- q^1)_{i,j} + (\partial_y^- q^2)_{i,j}$ and (3.3).

Many other dual formulations are possible. For example, if rather than relaxing the constraint $\|f - u\|^2 \leq \sigma^2$, we relax the equivalent one, $\|f - u\| \leq \sigma$, we obtain the Lagrangian version of (2.5),

$$(3.18) \quad \min \sum_{1 \leq i, j \leq n} \|\partial^+ u_{i,j}\| + \lambda \|f - u\|,$$

which can be reformulated as the following SOCP:

$$(3.19) \quad \begin{aligned} \min \quad & \sum_{1 \leq i, j \leq n} t_{i,j} + \lambda v_0 \\ \text{s.t.} \quad & (\partial_x^+ u)_{i,j} + (v_{i+1,j} - v_{i,j}) = f_{i+1,j} - f_{i,j} \quad \text{for } i = 1, \dots, n-1, j = 1, \dots, n, \\ & (\partial_y^+ u)_{i,j} + (v_{i,j+1} - v_{i,j}) = f_{i,j+1} - f_{i,j} \quad \text{for } i = 1, \dots, n, j = 1, \dots, n-1, \\ & (t_{i,j}; (\partial_x^+ u)_{i,j}, (\partial_y^+ u)_{i,j}) \in \mathcal{K}^3 \quad \text{for } i, j = 1, \dots, n-1, \\ & (t_{n,j}; (\partial_y^+ u)_{n,j}) \in \mathcal{K}^2 \quad \text{for } j = 1, \dots, n, \\ & (t_{i,n}; (\partial_x^+ u)_{i,n}) \in \mathcal{K}^2 \quad \text{for } i = 1, \dots, n, \\ & (v_0; v) \in \mathcal{K}^{n^2+1}. \end{aligned}$$

By assigning dual variables $r_{i,j}^1$ and $r_{i,j}^2$ to each of the first two constraints and associating $(s_{i,j}^0; s_{i,j}^1, s_{i,j}^2)$ to $(t_{i,j}; (\partial_x^+ u)_{i,j}, (\partial_y^+ u)_{i,j})$ and $(w_0; w)$ to $(v_0; v)$, we obtain the SOCP dual of (3.19):

$$(3.20) \quad \begin{aligned} \max \quad & \sum_{1 \leq i, j \leq n, i \neq n} r_{i,j}^1 (f_{i+1,j} - f_{i,j}) + \sum_{1 \leq i, j \leq n, j \neq n} r_{i,j}^2 (f_{i,j+1} - f_{i,j}) \\ \text{s.t.} \quad & -(\partial_x^- r^1)_{i,j} - (\partial_y^- r^2)_{i,j} + w_{i,j} = 0 \quad \text{for } i, j = 1, \dots, n, \\ & r_{i,j}^1 + s_{i,j}^1 = 0 \quad \text{for } i = 1, \dots, n-1, j = 1, \dots, n, \\ & r_{i,j}^2 + s_{i,j}^2 = 0 \quad \text{for } i = 1, \dots, n, j = 1, \dots, n-1, \\ & s_{i,j}^0 = 1 \quad \text{for } i, j = 1, \dots, n, \\ & w_0 = \lambda, \\ & (s_{i,j}^0; s_{i,j}^1, s_{i,j}^2) \in \mathcal{K}^3 \quad \text{for } i, j = 1, \dots, n-1, \\ & (s_{n,j}^0; s_{n,j}^2) \in \mathcal{K}^2 \quad \text{for } j = 1, \dots, n, \\ & (s_{i,n}^0; s_{i,n}^1) \in \mathcal{K}^2 \quad \text{for } i = 1, \dots, n, \\ & (w_0; w) \in \mathcal{K}^{n^2+1}. \end{aligned}$$

By eliminating variables $r_{i,j}^1$, $r_{i,j}^2$, $s_{i,j}^0$, and w_0 , this maximization problem can be simplified to the following minimization problem:

$$(3.21) \quad \begin{aligned} \min \quad & \sum_{1 \leq i, j \leq n, i \neq n} s_{i,j}^1 (f_{i+1,j} - f_{i,j}) + \sum_{1 \leq i, j \leq n, j \neq n} s_{i,j}^2 (f_{i,j+1} - f_{i,j}) \\ \text{s.t.} \quad & (\partial_x^- s^1)_{i,j} + (\partial_y^- s^2)_{i,j} + w_{i,j} = 0 \quad \text{for } i, j = 1, \dots, n, \\ & (1; s_{i,j}^1, s_{i,j}^2) \in \mathcal{K}^3 \quad \text{for } i, j = 1, \dots, n-1, \\ & (1; s_{n,j}^2) \in \mathcal{K}^2 \quad \text{for } j = 1, \dots, n, \\ & (1; s_{i,n}^1) \in \mathcal{K}^2 \quad \text{for } i = 1, \dots, n, \\ & (\lambda; w) \in \mathcal{K}^{n^2+1}. \end{aligned}$$

4. Nested dissection. SOCPs can be solved very efficiently in practice and, in theory, in polynomial time by interior-point methods. For example, path-following interior-point methods generate a sequence of interior points (i.e., points in the interior of \mathcal{K}) that follow the so-called central path toward an optimal solution. At each iteration, a set of primal-dual equations, which depend on current variable values and a duality gap parameter μ , are solved to yield an improving search direction. Rather than solving this system of equations directly, most implementations of interior-point methods first reduce these equations to a smaller system of linear equations involving a symmetric positive definite matrix that has the form AFA^\top . This system of linear equations is solved by computing the Cholesky factorization LDL^\top of AFA^\top , where L is a unit lower triangular matrix and D is a positive diagonal matrix. Although A remains fixed from one iteration to the next, F depends on the current primal-dual solution, so AFA^\top must be factorized at every iteration. Even though interior-point methods terminate in a polynomial number of iterations, each iteration can be very time consuming due to the $O(m^3)$ computational complexity of solving the Cholesky factorization of the $m \times m$ matrix AFA^\top , or even fail because of the lack of enough memory to store a dense $m \times m$ lower triangular matrix L . In our applications, m is quite large (e.g., in the SOCP (3.2), $m = 2n(n-1) + 1$, so that for $n = 1024$, m is roughly 2×10^6).

This difficulty can be overcome by exploiting the sparsity of the matrix A and the structure of F . In all interior-point methods, F is a block diagonal matrix $F \equiv \text{Diag}\{F_1, \dots, F_r\}$, where each of the diagonal submatrices $F_i \in \mathbb{R}^{n_i \times n_i}$ can be expressed as a sum of a scaled identity matrix and two rank-one terms (or for some methods, three rank-one terms), one with a positive coefficient and the other with a negative coefficient. Hence,

$$(4.1) \quad AFA^\top = \sum_{i=1}^r A_i F_i A_i^\top,$$

where $F_i = \gamma_i I + \mathbf{u}_i \mathbf{u}_i^\top - \mathbf{v}_i \mathbf{v}_i^\top$ (\mathbf{u}_i and \mathbf{v}_i are two columns vectors, not the signal and noise defined before).

It is easily verified that for the SOCP (3.2) the terms $A_i F_i A_i^\top$ have the same sparsity structure as $A_i A_i^\top$ for all i corresponding to two- and three-dimensional cones \mathcal{K}^2 and \mathcal{K}^3 . However, for the one large cone \mathcal{K}^{n^2+1} , $A_i F_i A_i^\top$ has a totally dense $(n^2 + 1) \times (n^2 + 1)$ block, whereas $A_i A_i^\top$ is very sparse. To take advantage of the sparsity of the latter matrix, implementations of interior-point methods compute the Cholesky factorization of only the sparse part of AFA^\top —i.e., for large cones, they include only the sparse term $\gamma_i A_i A_i^\top$ in the sum (4.1) when computing this factorization and then update the factorization to account for the dense part using either a *product-form* approach (see [11]) or a version of the Sherman–Morrison–Woodbury updating identity (see [2]). The latter approach is simpler but unfortunately less numerically stable than the former. The interior-point codes that we used in our experiments treated large cones in this fashion; hence in the following discussion we consider only the sparsity structure of $A_i A_i^\top$.

It is well known that sparsity in a symmetric matrix M does not guarantee that the Cholesky factor L will be sparse. For example, Cholesky factorization of the

rather sparse matrix

$$(4.2) \quad M = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & 2 & 0 & 0 & 0 \\ 1 & 0 & 3 & 0 & 0 \\ 1 & 0 & 0 & 7 & 0 \\ 1 & 0 & 0 & 0 & 43 \end{bmatrix}$$

yields $D = I$ and

$$(4.3) \quad L = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 1 & -1 & 1 & 0 & 0 \\ 1 & -1 & -2 & 1 & 0 \\ 1 & -1 & -2 & -6 & 1 \end{bmatrix}.$$

However, by reversing the order of the rows and columns of M , we obtain a matrix \tilde{M} whose Cholesky factorization yields $\tilde{D} = \text{Diag}(43, 7, 3, 2, 1)$ and

$$(4.4) \quad \tilde{L} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 1/43 & 1/7 & 1/3 & 1/2 & 1 \end{bmatrix}.$$

Nested dissection [10] is a scheme for symmetrically permuting the rows and columns of a symmetric matrix M so that the resulting Cholesky factorization takes optimal advantage of certain types of sparsity in M . To describe nested dissection, we let G_M be the *adjacency graph* associated with M . If $M \in \mathbb{R}^{m \times m}$, then G_M has m nodes and has an edge (i, j) if the (i, j) th element of M is nonzero. Nested dissection recursively applies the following *divide and conquer* strategy. Given a graph G_M , find a *separator* set of nodes N_S , whose removal from G_M divides it into two unconnected components G_U and G_V with corresponding node sets N_U and N_V . If the nodes in N_S in G_M are numbered after those in N_U and N_V , then the matrix M corresponding to this row and column ordering has the form

$$(4.5) \quad M = \begin{bmatrix} M_{UU} & \mathbf{0} & M_{US} \\ \mathbf{0} & M_{VV} & M_{VS} \\ M_{US}^\top & M_{VS}^\top & M_{SS} \end{bmatrix},$$

and most importantly, the Cholesky factor L of M inherits a similar structure, i.e.,

$$(4.6) \quad L = \begin{bmatrix} L_{UU} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & L_{VV} & \mathbf{0} \\ L_{SU} & L_{SV} & L_{SS} \end{bmatrix}.$$

This strategy is then applied recursively to the subgraphs G_U and G_V .

Let us now consider the application of nested dissection to the matrix AA^\top arising from the SOCP (3.2), which as explained above corresponds to the sparse part of AFA^\top . $(AA^\top)_{l,k} \neq 0$ if and only if the l th and k th constraints share decision variables. The equality constraints in our problem are

$$(4.7) \quad (\partial_x^+ u)_{i,j} + (v_{i+1,j} - v_{i,j}) = f_{i+1,j} - f_{i,j} \quad \text{for } i = 1, \dots, n-1, j = 1, \dots, n,$$




FIG. 1. Constraint dependency of (3.2).




FIG. 2. Constraint dependency of (3.17).

$$(4.8) \quad (\partial_y^+ u)_{i,j} + (v_{i,j+1} - v_{i,j}) = f_{i,j+1} - f_{i,j} \quad \text{for } i = 1, \dots, n, j = 1, \dots, n-1,$$

$$(4.9) \quad v_0 = \sigma.$$

Clearly, the last constraint shares no variable with any other constraint; thus its position in the ordering of the constraints is irrelevant. Let $C_{i,j}^1$ and $C_{i,j}^2$ denote the (i,j) th constraints in (4.7) and (4.8), respectively. For a particular pair (i,j) (i.e., row of A), $C_{i,j}^1$ shares at least one decision variable with the following constraints, given their existence:

$$\begin{array}{lll} C_{i-1,j}^1, & C_{i,j-1}^2, & C_{i,j}^2 \quad (\text{sharing } v_{i,j}), \\ C_{i+1,j}^1, & C_{i+1,j-1}^2, & C_{i+1,j}^2 \quad (\text{sharing } v_{i+1,j}). \end{array}$$

$C_{i,j}^2$ shares at least one decision variable with the following constraints, given their existence:

$$\begin{array}{lll} C_{i-1,j}^1, & C_{i,j}^1, & C_{i,j-1}^2 \quad (\text{sharing } v_{i,j}), \\ C_{i-1,j+1}^1, & C_{i,j+1}^1, & C_{i,j+1}^2 \quad (\text{sharing } v_{i,j+1}). \end{array}$$

The constraint dependency relations (i.e., adjacency graph G_M of AA^\top excluding the row and column corresponding to (4.9), which we shall refer to as M) are depicted in Figure 1 for the case of $n = 8$, in which squares represent $C_{i,j}^1$ and bullets $C_{i,j}^2$. Clearly, removal of the row of squares $C_{4,j}^1$, $j = 1, \dots, 8$, separates the dependency graph in Figure 1 into two independent parts. If we let $S \equiv \{C_{4,j}^1 : j = 1, \dots, 8\}$ and let U and V denote the sets of constraints corresponding to the upper and the lower independent parts, and order the constraints in U before those in V and those in V before those in S , then M has the form of (4.5).

Furthermore, removal of the column of bullets $C_{i,4}^2$, $i = 1, \dots, 4$, separates U into two independent parts— U_L and U_R —and removal of the column of bullets $C_{i,4}^2$, $i = 5, \dots, 8$, separates V into two independent parts— V_L and V_R . As a result, M_{UU} and M_{VV} can be ordered so that each has the form (4.5). This process is then repeated recursively until every node has been assigned to some separator set (i.e., every row and column of M has been numbered). In the appendix, we analyze the storage cost and the number of multiplications for factorizing $M = AA^\top$ corresponding to the

equality constraints (4.7) and (4.8) after ordering this matrix by nested dissection. These costs are $(31/4)n^2 \log_2 n + O(n^2)$ and $943n^3/84 + O(n^2 \log_2 n)$ for an $n \times n$ image, respectively.

Finally, we depict in Figure 2 the dependency graph of the dual SOCP formulation (3.17). In Figure 2 the (i, j) th bullet $\bar{C}_{i,j}^1$ represents the constraint $(\operatorname{div} q)_{i,j} - w_{i,j} = f_{i,j}$, and the (i, j) th square $\bar{C}_{i,j}^2$ represents the constraint $q_{i,j}^0 = 1$. Since constraints $\bar{C}_{i,j}^1$, $\bar{C}_{i+1,j}^1$, and $\bar{C}_{i,j+1}^1$ contain either $q_{i,j}^1$ or $q_{i,j}^2$ or both of them, and the 3×3 fully dense block F_k corresponding to cone $(q_{i,j}^0; q_{i,j}^1; q_{i,j}^2)$ causes the 3 corresponding columns in A to merge, constraint $\bar{C}_{i,j}^2$ in $AF A^T$ essentially shares variables with constraints $\bar{C}_{i,j}^1$, $\bar{C}_{i+1,j}^1$, and $\bar{C}_{i,j+1}^1$; hence, in Figure 2, $\bar{C}_{i,j}^2$ is connected to $\bar{C}_{i,j}^1$, $\bar{C}_{i+1,j}^1$, and $\bar{C}_{i,j+1}^1$. The connections between bullets $\bar{C}_{i,j}^1$'s are due to the sharing of $q_{i,j}^1$ or $q_{i,j}^2$ for some i, j .

5. Numerical results. In this section, we present numerical results that illustrate the effectiveness of our SOCP approach and compare it to the PDE approach. We then discuss the use of domain decomposition to reduce the storage and computational effort needed by the SOCP approach.

5.1. SOCP/PDE comparison. We used the commercial optimization package Mosek as our SOCP solver. Mosek is designed to solve a variety of large-scale optimization problems, including SOCPs. Before solving a large-scale problem, Mosek uses a presolver to remove redundant constraints and variables and to reorder constraints. The Mosek presolver reorders the coefficient matrix A using a combination of a graph partitioning-based method and an approximate minimum local-fill-in ordering method to increase the sparsity of the Cholesky factors of $AF A^T$. This general-purpose reordering produced results for our problem that were very good, but not as good as nested dissection, which further reduced the time per iteration and the total time by approximately 30 percent in our tests on a 512×512 image. Therefore, we turned off Mosek's reordering heuristic and replaced it by our own nested dissection reordering code in our tests.

Existing PDE approaches require the ROF model to be represented in the relaxed form (1.2) and the existence of the Euler-Lagrange equation $g(u) = 0$, where $g(u)$ is given by (1.3). To handle the case of $\|\nabla u\| = 0$, most PDE approaches add a small perturbation $\epsilon > 0$ to $\|\nabla u\|$ in (1.3). The PDE method that we used in our numerical tests is an artificial time stepping method which computes the steady state solution u of the parabolic system

$$(5.1) \quad \frac{\partial u}{\partial t} = \frac{1}{2\lambda} \nabla \cdot \left(\frac{\nabla u}{\|\nabla u\| + \epsilon} \right) + f - u.$$

To compare the two approaches for each noisy image, we first solved the SOCP formulation (3.2) of (1.1) and calculated λ from the optimal dual solution using the following theorem.

THEOREM 5.1. *The discrete versions of problems (1.1) and (1.2) have the same optimal solution u if $2\lambda\sigma = z_0$, where z_0 is the first entry of $(z_0, z) \in \mathcal{K}^{n^2+1}$ and (z_0, z) is the dual vector corresponding to (v_0, v) in the SOCP formulation (3.2) of problem (1.1). (Recall that $\int_{\Omega} |v|^2 \, dx \leq \sigma^2$ is represented in the discrete form by $(v_0, v) \in \mathcal{K}^{n^2+1}$ with $v_0 = \sigma$.)*

We then used this Lagrange multiplier λ in the PDE method. In practice, it is much easier to estimate the noise level of an image than it is to estimate an appropriate

TABLE 1
Numerical comparison between SOCP (MATLAB+Mosek) and PDE (C++) approaches.

| Name | Size $n \times n$ | Fig. | SOCP result u_1 | vs. PDE result u_2 |
|---------|----------------------|-------|--------------------------|---------------------------|
| | | | $\max_{i,j} u_1 - u_2 $ | $\ u_1 - u_2\ _{L^1}/n^2$ |
| Books | 64×64 | 3–5 | 0.58 | 0.04 |
| Dotplus | 100×100 | 6–8 | 0.97 | 0.10 |
| Scale | 252×252 | 9–11 | 0.74 | 0.02 |
| Barbara | 256×256 | 12–14 | 0.90 | 0.05 |
| Barbara | 350×350 | 15–17 | 0.95 | 0.06 |
| Barbara | 512×512 | 24–26 | 0.96 | 0.06 |

λ . Experiments show that a σ that is a little smaller than the true one (i.e., a slight amount of undersmoothing) yields the best results.

In our tests, we called Mosek from MATLAB, while we called the PDE-based code (an improved version of a C++ code provided to us by Stanley Osher and Jinjun Xu of UCLA) directly. The running times given in Table 2 for the problems whose sizes are given in Table 1 are the total times in seconds needed to solve the problem on a Sun E450 with four 300MHz SuperSparc II processors and 4GB memory. Since both codes are single-threaded, they run only on one processor and use 1GB memory at a time. The SOCP run times do not include the time taken by our nested dissection code to reorder constraints and the time taken by Mosek for presolving, since the process of constraint reordering and presolving remains the same for all images of the same size.

During their execution, both approaches record the objective values of (1.2):

$$(5.2) \quad z(t) = \int_{\Omega} \|\nabla u_t\| \, dx + \lambda \int_{\Omega} |f - u_t|^2 \, dx.$$

The objective values obtained by the PDE approach decrease with time t over the long run but fluctuate and never reach as low a value as that obtained by the SOCP approach. Given a nearly optimal objective value, the PDE approach may have to iterate further for a large number of steps to obtain a lower objective value. For these reasons, we used the following stopping rules for the PDE approach:

- initially, let $z_{\min} = \infty$ and $t_{\min} = 0$, where z_{\min} and t_{\min} denote the minimal objective value and its time;
- at time t , update $z_{\min} := z(t)$ and $t_{\min} := t$ if $z_{\min} < z(t) - \lambda$;
- to denoise an image of size $n \times n$, stop at time t once $t - t_{\min} \geq \max\{0.15t_{\min}, n\}$.

These stopping criteria ensure that the PDE method spends a reasonable amount of time searching for a possibly lower objective value but does not waste too much time trying to obtain an insignificant improvement.

In all of our tests starting from the same noisy image, the two approaches produced images that were visually indistinguishable from one another. Table 1 shows that in all our tests the gray values (u_1 and u_2) produced by the two approaches differ by at most 0.97 over all pixels with an average difference of less than 0.10. Therefore, we present in Figures 3–17, in addition to the original and noisy images, only the SOCP images. Similar results were obtained on the other images.

From Table 2, one can see that the SOCP approach requires a smaller number of iterations than the PDE approach while every iteration takes more time. Note that as the size of the image grows, the number of iterations required by the SOCP approach remains relatively constant, while that number grows for the PDE approach. In spite




FIG. 3. *Original books*
 64×64 .




FIG. 4. *Noisy books*
($\sigma = 20$).




FIG. 5. *Denoised image by the SOCP method.*




FIG. 6. *Original dotplus*
 100×100 .




FIG. 7. *Noisy dotplus*
($\sigma = 40$).




FIG. 8. *Denoised image by the SOCP method.*




FIG. 9. *Original scale*
 252×252 .




FIG. 10. *Noisy scale*
($\sigma = 20$).




FIG. 11. *Denoised image by the SOCP method.*



FIG. 12. *Original Barbara*
 256×256 .



FIG. 13. *Noisy Barbara*
($\sigma = 20$).



FIG. 14. *Denoised image by the SOCP method.*

of this, the SOCP approach took roughly between one and two times the total time that the PDE approach took. This is because, in theory, the PDE method takes $O(n^2)$ CPU time per iteration for an $n \times n$ image, while the SOCP method (implemented


FIG. 15. *Original Barbara* 350×350 .FIG. 16. *Noisy Barbara* ($\sigma = 20$).FIG. 17. *Denoised image by the SOCP method.*

TABLE 2
Numerical comparison between SOCP (MATLAB+Mosek) and PDE (C++) approaches (cont.).

| Name | λ | SOCP | | | PDE ($\Delta t = 0.1$) | | |
|---------|--------------|------------|------|------------|--------------------------|------|------------|
| | | Total time | Itr. | Obj. value | Total time | Itr. | Obj. value |
| Books | 0.0374718930 | 4.42 | 15 | 128927 | 3.22 | 950 | 128976 |
| Dotplus | 0.0180513895 | 13.68 | 16 | 592183 | 14.90 | 1875 | 592438 |
| Scale | 0.0337088230 | 157.86 | 17 | 3706529 | 81.21 | 1479 | 3706980 |
| Barbara | 0.0374657886 | 130.13 | 16 | 1691196 | 61.73 | 1166 | 1692180 |
| Barbara | 0.0374969349 | 283.56 | 16 | 2905191 | 151.71 | 1464 | 2907500 |
| Barbara | 0.0375681804 | 806.70 | 16 | 5895359 | 456.25 | 1900 | 5900690 |

as described in section 4) takes $O(n^3)$ CPU time per iteration. In our numerical tests the SOCP approach took between $O(n^2 \log n)$ and $O(n^3)$ time per iteration due to the effect of memory references, which are slower than floating-point operations.

We also measured the running times on a PC with a Pentium IV 1.7GHz CPU and 512MB memory. We used the MS Windows versions of MATLAB, Mosek, and GUN g++ compiler. On this newer computer, the SOCP code ran 2–2.5 times faster and the PDE code ran about 4 times faster solving the same set of problems. Due to the PC's smaller amount of memory and cache, the SOCP approach, which performs more memory fetches, got a smaller speedup than the PDE approach, exhibiting a

running time somewhere between the $n \times n$ -image storage cost of $O(n^2)$ and the nested dissection Cholesky factorization storage cost of $O(n^2 \log n)$.

5.2. Domain decomposition. Even with nested dissection the SOCP approach requires $O(n^2 \log n)$ memory. Hence, to denoise large (or three-dimensional) images, we consider here a *domain decomposition* approach. Let $\tilde{n} \approx n/k$, where k is a positive integer and $\tilde{n} > n/k$. By subdividing an $n \times n$ image into $k^2 \tilde{n} \times \tilde{n}$ subimages with some overlap and applying the SOCP method to each of the subimages, we can reduce the storage needed by a factor of approximately $1/k^2$. Since, in applying the SOCP method to each subimage, the number of multiplications per iteration is reduced by a factor of approximately $1/k^3$ compared with applying the SOCP method to the entire image, and there are k^2 subimages, the total CPU time should also be reduced by a factor of approximately $1/k$. After obtaining the denoised subimages, we discarded the values of u for the rows (columns) of pixels in that half of any region of overlap closest to the boundary of the subimage. For example, Figure 20 illustrates a 2×2 domain decomposition with a 4-pixel overlap between subimages. After obtaining the denoised 65×65 image u for the top left block, we discarded the u values for the rightmost two columns and bottommost two rows of pixels. These values have indices (i, j) with $i = 64, 65$ and $j = 1, \dots, 65$, and $j = 64, 65$ and $i = 1, \dots, 65$. An analogous statement applies to the overlapped pixels of the other 3 denoised subimages.

Due to the different graphical characteristics of the subimages, including signal-to-noise ratio, contrast, and homogeneity, applying the SOCP method to all subimages using the same σ does not usually give the same output as applying the SOCP method to the entire image. In some cases the results obtained by using the former are better, while in other cases they are worse. The following numerical example illustrates this. Figure 18 is the noisy image obtained by selecting a rectangular area of the noisy image in Figure 13 and removing the outside portion of the image. Figure 19 is the result of the ROF method (3.2) separately applied to 4 equal subimages of size 65×65 with the same $\sigma = 18 \times 65^2$. In the numerical tests, we chose the size of all subimages so that all adjacent image pairs had overlapping regions that were 4 pixels wide. Figure 20 illustrates the 2×2 domain decomposition applied to the image in Figure 18. Figure 21 is the result of ROF applied to the entire image with $\sigma = 18 \times 65^2$. In Figure 21, Barbara's left eye is clearer than in Figure 19. This shows that the subimage at the bottom right is oversmoothed. After reducing $\sigma_{2,2}$ from 18×65^2 to 16×65^2 for this oversmoothed subimage, the ROF method yields Figure 22, in which Barbara's left eye is much clearer. Figure 23 gives the ten-times amplified differences between Figures 21 and 22. The differences in the regions of the four subimages vary, but we cannot see any big differences by comparing Figures 21 and 22. However, as we shall now demonstrate, we can obtain even closer results.

In [20], Strong and Chan give exact analytical solutions to the relaxed TV denoising problem (1.2) applied to one-, two-, and three-dimensional piecewise constant images and smooth radially symmetric ones. Their results establish an exact relationship between the Lagrange multiplier λ and the effect of TV regularization (1.2) on an image feature and show its localness property (i.e., in a piecewise constant or smooth radially symmetric region the solution of (1.2) depends only on λ and whether or not the region is on the image boundary, and is independent to the features of other regions). This can be extended approximately to complicated noisy images since they are composed of many tiny piecewise constant or smooth radially symmetric regions. Therefore, we can expect solving (1.2) with the same λ (or σ 's for all subimages so that $z_0/(2\sigma) = \lambda$) as used for the entire image to give close results. Since TV



FIG. 18. *Noisy image* ($\sigma = 20 \times 65^2$).



FIG. 19. *2 x 2 partitioned result* ($\sigma_{1,1} = \sigma_{1,2} = \sigma_{2,1} = \sigma_{2,2} = 18 \times 65^2$).




FIG. 20. *2 x 2 domain decomposition applied to image of size 126 x 126.*



FIG. 21. *Nonpartition result* ($\sigma = 18 \times 65^2$).



FIG. 22. *2 x 2 partitioned result* ($\sigma_{1,1} = \sigma_{1,2} = \sigma_{2,1} = \sigma_{2,2} = 16 \times 65^2$).



FIG. 23. *Ten-times amplified difference between Figures 21 and 22.*

regulation treats regions on boundaries differently, domain decomposition may give different results near subimage boundaries. Sharp edges are different from sharp noise by the surrounding geometry, e.g., two sides of sharp edges have very different brightness. Therefore, if a sharp edge is very close to the image boundaries, the difference between edges and noise is weaker. Moreover, every pixel in an image may affect the denoising results of the whole image, and their dependency is ambiguous. This effect, however, is limited to strips a few pixels wide when noise is present. Hence, in our tests, we discarded redundant information in these boundary strips of the subimages.

In our first test of the above approach we applied a 4×4 domain decomposition to a 512×512 image. The results are depicted in Figures 24–26. All subimages were of size 137×137 , giving overlapping regions between adjacent subimages that were 12 pixels wide. The image values for the 6 rows (columns) of pixels in the overlapping regions closest to the boundaries of each subimage were discarded. The σ 's for the full image and all of the subimages were chosen so that $z_0/(2\sigma) = 0.0375$. Table 3 gives the numerical results. The results obtained by the two approaches by more than 1 on the 256 gray-level scale was 0.11% of the total 512×512 pixels. The average difference per pixel was 0.07 and the maximum difference was +0.97 and -1.15 . The running time 1280 seconds for denoising the entire image was estimated based on the $O(n^3)$ theoretical CPU time. The difference was more than 1 on the 256 gray-level scale for only 0.0008% of the pixels. The average difference per pixel was 0.04 and the maximum difference was +0.75 and -1.31 . Clearly, SOCP with domain decomposition gave nearly identical results to SOCP applied to the full image.

FIG. 24. *Original 512 × 512 image.*FIG. 25. *Noisy 512 × 512 image ($\sigma = 20$).*FIG. 26. *4 × 4 domain decomposition applied to Figure 25.*

TABLE 3

4 × 4 domain decomposition applied to Figure 24. *Blk.* = block (subimage) index; *itr.* = number of SOCP iterations for denoising subimage; *time* = SOCP running time in seconds for denoising subimage.

| Blk. | Itr. | Time | Blk. | Itr. | Time | Blk. | Itr. | Time | Blk. | Itr. | Time |
|--|------|-------|-------|------|-------|-------|------|-------|-------|--------|-------|
| (1,1) | 19 | 33.03 | (1,2) | 20 | 34.84 | (1,3) | 18 | 32.01 | (1,4) | 17 | 29.95 |
| (2,1) | 18 | 31.97 | (2,2) | 18 | 31.74 | (2,3) | 17 | 29.73 | (2,4) | 17 | 30.25 |
| (3,1) | 18 | 31.34 | (3,2) | 17 | 30.61 | (3,3) | 20 | 35.70 | (3,4) | 18 | 31.18 |
| (4,1) | 17 | 30.10 | (4,2) | 17 | 31.06 | (4,3) | 18 | 32.03 | (4,4) | 18 | 32.72 |
| Sum of the total time (seconds) for denoising all subimages (Figure 26): | | | | | | | | | | 508.26 | |
| Total time for denoising the entire image: | | | | | | | | | | 806.70 | |

In this example the total computation time for the domain decomposition approach was about 37 percent less than the time required by the SOCP method applied to the full image. Since $16 \cdot 137^3 / 512^3 \approx 0.33$, one would expect something close to a 67 percent reduction in running time. The reason this was not the case is explained by our remark in section 5.1 about execution time and memory references when cache is limited and slow.

We also applied an 8×8 domain decomposition to the noisy 512×512 image in

TABLE 4

8×8 domain decomposition applied to Figure 24. Time (itr.) = SOCP running time in seconds (iterations) for denoising subimage

| Block index | | Column | | | | | | | |
|--|---|----------|----------|----------|----------|----------|----------|----------|----------|
| | | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
| Row | 1 | 6.86(16) | 6.56(16) | 8.37(19) | 6.24(15) | 5.80(14) | 5.32(13) | 5.67(14) | 6.12(14) |
| | 2 | 7.07(17) | 6.90(16) | 6.51(16) | 5.70(14) | 6.04(14) | 6.02(14) | 6.38(15) | 5.83(14) |
| | 3 | 5.65(14) | 6.64(16) | 6.90(16) | 7.53(17) | 7.02(16) | 5.86(14) | 5.77(14) | 6.02(15) |
| | 4 | 6.76(16) | 7.17(17) | 7.21(17) | 5.79(14) | 5.88(13) | 5.84(14) | 5.38(13) | 6.29(15) |
| | 5 | 6.19(15) | 7.19(17) | 6.63(16) | 6.17(14) | 5.93(14) | 5.75(14) | 6.74(16) | 6.97(17) |
| | 6 | 5.99(14) | 6.75(16) | 6.88(17) | 7.06(16) | 6.80(15) | 5.73(14) | 6.50(15) | 6.21(15) |
| | 7 | 7.06(17) | 7.22(17) | 7.16(16) | 6.59(15) | 6.35(15) | 6.35(15) | 6.71(16) | 6.14(15) |
| | 8 | 7.04(17) | 7.58(17) | 7.10(17) | 7.83(19) | 7.25(17) | 6.39(15) | 7.13(17) | 7.57(18) |
| Sum of the total time (seconds) for denoising all subimages: | | | | | | | | | 418.06 |

Figure 25 after extending it to one with a size of 516×516 so that all subimages were of the same size (75×75) and the overlapping regions remained 12 pixels wide. As in the previous test, the σ 's were chosen so that $z_0/(2\sigma) = 0.0375$ for all subimages. The numerical results given in Table 4 show that the reduction in the total computation time was further improved from 37 percent to 52 percent. Moreover, the total run time of 418.06 seconds was less than the 456.25-second run time of the PDE approach. The difference between the results obtained by applying the SOCP method with domain decomposition and applying it to the entire image was more than 1 on the 256 gray-level scale for 0.0015% of the pixels. The average difference per pixel increased to 0.08 and the maximum difference was +1.47 and -1.06. This difference was still not perceivable to the naked eye.

6. SOCP formulations of other restoration models. Meyer [15] introduced the following image restoration model based on TV and the space of oscillating functions:

$$(6.1) \quad \min \int_{\Omega} \|\nabla u\| dx + \lambda \|v\|_* \quad \text{s.t. } f = u + v,$$

where $\|v\|_*$ is defined as the infimum of $\|\sqrt{g_1^2(x, y) + g_2^2(x, y)}\|_{L^\infty}$ over all $g_1, g_2 \in L^\infty(\mathbb{R}^2)$ satisfying $v = \partial_x g_1 + \partial_y g_2$. To handle $\|\sqrt{g_1^2(i, j) + g_2^2(i, j)}\|_{L^\infty}$ in the discretized version of (6.1), we introduce a three-dimensional quadratic cone

$$(g_0(i, j); g_1(i, j), g_2(i, j)) \in \mathcal{K}^3$$

for each i, j and let $w \geq g_0(i, j)$ for all i, j . The infimum of $\|\sqrt{g_1^2(i, j) + g_2^2(i, j)}\|_{L^\infty}$ is then equal to the minimum of w . Thus we can solve a discretized problem of (6.1) as an SOCP. It is important to note that because Meyer's model (6.1) is not amenable to the PDE approach, it has been primarily of theoretical interest. Thus, the SOCP approach can be applied to image restoration models that cannot be solved by a PDE approach.

In [22], Vese and Osher proposed an approximate version of Meyer's model and showed that their PDE-based method can decompose a given image f into u and v , where u and v represent cartoon and texture, respectively. Their method solves the minimization problem

$$(6.2) \quad \min \int_{\Omega} \|\nabla u\| dx + \lambda \int |f - u - \partial_x g_1 - \partial_y g_2| + \mu \left[\int (\sqrt{g_1^2 + g_2^2})^p \right]^{1/p},$$

where λ and μ are tuning parameters and $p \geq 1$. Inequalities with p -norm for p being rational can be represented as a system of second-order conic inequalities (refer to [1]). Therefore, (6.2), after discretization, can be formulated as an SOCP and solved by interior-point methods. In a forthcoming paper [24], we solve Meyer's model as an SOCP and compare the solutions with those obtained by the Vese–Osher approximation.

In [14], Lysaker, Lundervold, and Tai introduced a method for image restoration based on solving the following two-dimensional problem:

$$(6.3) \quad \begin{aligned} \min \quad & \int_{\Omega} |u_{x_1 x_1}| + |u_{x_2 x_2}| \, dx \\ \text{s.t.} \quad & u + v = f, \\ & \int_{\Omega} |v|^2 \, dx_1 \, dx_2 \leq \sigma^2. \end{aligned}$$

The discretization of the second derivative of $u = \{u_{i,j}\}$ is linear in the $u_{i,j}$'s. To handle absolute values in minimization problems, say $\min |z|$, one can introduce an extra variable and transform $\min |z|$ into equivalent problems:

$$(6.4) \quad \min |z|$$

$$(6.5) \quad \iff \min t \quad \text{s.t.} \quad z \leq t, \quad -z \leq t$$

$$(6.6) \quad (s \stackrel{\text{def}}{=} t + z) \iff \min(s - z) \quad \text{s.t.} \quad 2z \leq s, \quad s \geq 0.$$

Both (6.5) and (6.6) have a linear objective subject to linear constraints. In some cases, (6.6) is preferred since $s \geq 0$ is easier to handle. Therefore, expressing $\int_{\Omega} |v|^2 \, dx \leq \sigma^2$ as a second-order conic constraint as before, it is clear that problem (6.3) can be formulated as an SOCP. This also works if the objective function is replaced by $\int_{\Omega} |u_{x_1 x_1}| + 2|u_{x_1 x_2}| + |u_{x_2 x_2}| \, dx$.

A variant of (1.1) was proposed by Osher, Solé, and Vese (OSV) in [16]. The model assumes that $f - u = \text{div} g$, where $g = (g_1, g_2)$, and there exists a unique decomposition $g = \nabla P + Q$, where Q satisfies $\text{div} Q = 0$. Therefore, $f - u = \text{div} g = \Delta P$ or $\Delta^{-1}(f - u) = P$. The OSV model is

$$(6.7) \quad \begin{aligned} \min \quad & \int_{\Omega} \|\nabla u\| \, dx \\ \text{s.t.} \quad & \int_{\Omega} \|\nabla(\Delta^{-1}(f - u))\| \, dx \leq \sigma^2, \end{aligned}$$

or equivalently

$$(6.8) \quad \begin{aligned} \min \quad & \int_{\Omega} \|\nabla u\| \, dx \\ \text{s.t.} \quad & \int_{\Omega} \|\nabla P\| \, dx \leq \sigma^2, \\ & f - u = \Delta P. \end{aligned}$$

Clearly, by using a finite difference form of the Laplace operator and expressing the objective function and the constraint $\int_{\Omega} \|\nabla P\| \, dx \leq \sigma^2$ as we have previously described, we can reformulate (6.8) as an SOCP.

In [8], the authors analyzed the ROF model with an L^1 fidelity term $\int_{\Omega} |v| \leq \sigma$. Clearly, by discretizing $\int_{\Omega} |v|$ as the sum $\sum_{i,j} |v_{i,j}|$ and handling the absolute values $|v_{i,j}|$ using (6.5) or (6.6), this problem can be formulated as an SOCP.

It is not difficult to see that various TV-based deblurring models can also be formulated as SOCPs.

Appendix. In this appendix, we present a detailed analysis of Cholesky factorization of a matrix M whose rows and columns have been symmetrically reordered by applying nested dissection to the adjacency graph of M , assuming that the latter is an $n \times n$ grid graph of the type shown in Figure 1. We follow an approach similar to the one given in [18]. First, let us introduce some graph notation.

A graph G is defined as the pair (V_G, E_G) , where V_G is the set of nodes and E_G is the set of edges of G . In Figure 27, each square or bullet is a node, and each line segment connecting a pair of nodes is an edge. Figure 27 has 8 rows and 7 columns corresponding to the C^2 constraints, and 7 rows and 8 columns corresponding to the C^1 constraints. For convenience, we number node $C_{i,j}^k$ by $(i, j)^k$. In this way, $C_{3,1}^1$ corresponds to node $(3, 1)^1$ and $C_{8,1}^2$ to node $(8, 1)^2$. Due to the boundary conditions, the row of nodes $(8, \cdot)^1$ and the column of nodes $(\cdot, 8)^2$, corresponding to $(\partial_x^+ u)_{8,j} = 0$ and $(\partial_y^+ u)_{i,8} = 0$ for $1 \leq i, j \leq 8$, do not appear in Figure 27. Let us denote by $G_{2n,2n}$ a graph of the form of Figure 27 with nodes $(i, j)^k$ for $k = 1, 2$ and $i, j = 1, \dots, n-1$, $(i, n)^1$ for $i = 1, \dots, n$, and $(n, j)^2$ for $j = 1, \dots, n$. For example, Figure 27 depicts a $G_{16,16}$ mesh.

Given the graph $G_{2n,2n}$, we define three extended graphs $A_{2n,2n}$, $B_{2n,2n}$, and $C_{2n,2n}$ as follows:

$$(A.1) \quad A_{2n,2n} \stackrel{\text{def}}{=} (V_A, E_A), \quad \text{where } V_A = V_G \bigcup L_{\text{left}} \bigcup L_{\text{bottom}},$$

$$(A.2) \quad B_{2n,2n} \stackrel{\text{def}}{=} (V_B, E_B), \quad \text{where } V_B = V_G \bigcup L_{\text{left}} \bigcup L_{\text{bottom}} \bigcup L_{\text{right}},$$

$$(A.3) \quad C_{2n,2n} \stackrel{\text{def}}{=} (V_C, E_C), \quad \text{where } V_C = V_G \bigcup L_{\text{top}} \bigcup L_{\text{left}} \bigcup L_{\text{bottom}} \bigcup L_{\text{right}},$$

where

$$(A.4) \quad L_{\text{left}} = \{(i, 0)^2 \mid i = 1, 2, \dots, n\},$$





$$(A.5) \quad L_{\text{right}} = \{(i, n)^2 \mid i = 1, 2, \dots, n\},$$

$$(A.6) \quad L_{\text{top}} = \{(0, j)^1 \mid j = 1, 2, \dots, n\},$$

$$(A.7) \quad L_{\text{bottom}} = \{(n, j)^1 \mid j = 1, 2, \dots, n\},$$

and E_A , E_B , and E_C are the sets of edges that connect V_A , V_B , and V_C by a grid of the type shown in Figure 27. For example, for $n = 8$, $A_{16,16}$, $B_{16,16}$, and $C_{16,16}$ are depicted in Figures 28, 29, and 30, in which extended boundaries are distinguished by being enclosed in ellipses.

The reason that we need these extended graphs in our analysis can be illustrated as follows. As we described in section 4, the graph $G_{16,16}$ can be separated into four subgraphs, each of which is a $G_{8,8}$, by separators $(4, j)^1$, $j = 1, \dots, 8$, followed by $(i, 4)^2$, $i = 5, \dots, 8$, and $(i, 4)^2$, $i = 1, \dots, 4$. The order of the graph separators is just the reverse of the order assigned to the rows and columns in the matrix. It is well known that, in Cholesky factorization (CF), eliminating a node of degree d (i.e., the number of nodes directly connected to it) costs $K_m = d(d+3)/2$ multiplies and requires $K_s = d+1$ storage. In the next stage the case is different from that of $G_{8,8}$

FIG. 27. $G_{16,16}$.FIG. 28. $A_{16,16}$.FIG. 29. $B_{16,16}$.FIG. 30. $C_{16,16}$.

since, for example, the subgraph $G_{8,8}$ in the upper right corner of $G_{16,16}$ is connected to the nodes $(4, j)^1$, $j = 5, \dots, 8$, and $(i, 4)^2$, $i = 1, \dots, 4$, which are the bottom and left extended boundaries of $G_{8,8}$. Therefore, we need to consider separating $A_{8,8}$ instead of $G_{8,8}$. Since applying any rotation or reflection to $A_{8,8}$ does not change the multiplication and storage costs of CF, separating any of the other three subgraphs of $G_{16,16}$ is essentially the same as separating $A_{8,8}$. Similar remarks apply to any rotation or reflection of $A_{2n,2n}$, $B_{2n,2n}$, and $C_{2n,2n}$.

Let $\mu(\cdot)$ denote a function of a graph whose value gives the complexity of applying CF to a matrix corresponding to that graph. Specifically, we use $\mu_m(\cdot)$ and $\mu_s(\cdot)$ to present the multiplication cost and the storage cost, respectively. In the divide and conquer approach, the complexity of applying CF to $G_{2n,2n}$ can be calculated by the recursion

$$(A.8) \quad \mu(G_{2n,2n}) = 4\mu(A_{n,n}) + \mu(S_{2n,2n}^G),$$

where $S_{2n,2n}^G$ is defined as the separator of $G_{2n,2n}$.

Following similar arguments, one can divide each of $A_{2n,2n}$, $B_{2n,2n}$, and $C_{2n,2n}$ into 4 subgraphs and derive the following recursions:

$$(A.9) \quad \mu(A_{2n,2n}) = \mu(A_{n,n}) + 2\mu(B_{n,n}) + \mu(C_{n,n}) + \mu(S_{2n,2n}^A),$$

$$(A.10) \quad \mu(B_{2n,2n}) = 2\mu(B_{n,n}) + 2\mu(C_{n,n}) + \mu(S_{2n,2n}^B),$$

$$(A.11) \quad \mu(C_{2n,2n}) = 4\mu(C_{n,n}) + \mu(S_{2n,2n}^C).$$

The next step is to calculate $\mu(S_{2n,2n}^G)$, $\mu(S_{2n,2n}^A)$, $\mu(S_{2n,2n}^B)$, and $\mu(S_{2n,2n}^C)$. Take $\mu(S_{16,16}^G)$, for example. Recall that we eliminate $(i, 4)^2$, $i = 1, \dots, 4$, followed by $(i, 4)^2$, $i = 5, \dots, 8$, and $(4, j)^1$, $j = 1, \dots, 8$. At the time we eliminate $(1, 4)^2$, all nodes except for those in $S_{2n,2n}^G$ are eliminated. Therefore, $(1, 4)^2$ is connected to $(i, 4)^2$, $i = 2, \dots, 4$, and $(4, j)^1$, $j = 1, \dots, 8$. Therefore, $d_{(1,4)^2} = 3 \times (n/2) - 1$, where $n = 8$. Next, $(2, 4)^2$ is eliminated with $d_{(2,4)^2} = 3 \times (8/2) - 2$. When we eliminate $(4, 4)^2$, it has degree of $3 \times (8/2) - 4 = 8$ since it is connected to eight nodes $(4, j)^1$, $j = 1, \dots, 8$. Hence, the sums of d and d^2 for $G_{2n,2n}$'s separator $(i, n/2)^2$, $i = 1, \dots, n/2$, are given by $\sum_{j=n}^{3n/2-1} j$ and $\sum_{j=n}^{3n/2-1} j^2$, respectively. By symmetry, this is also true for the separator $(i, n/2)^2$ for $i = n/2 + 1, n/2 + 2, \dots, n$. Last, we eliminate the n nodes $(n/2, j)$ for $j = 1, 2, \dots, n$. Whenever we eliminate one of them, it is connected to all the nodes to be eliminated. Therefore, for these last n nodes, the sums of d and d^2 are given by $\sum_{j=0}^{n-1} j$ and $\sum_{j=0}^{n-1} j^2$, respectively. Altogether,

$$(A.12) \quad \sum_{v \in S_{2n,2n}^G} d_v^2 = 2 \sum_{j=n}^{3n/2-1} j^2 + \sum_{j=0}^{n-1} j^2 = \frac{23}{12}n^3 + O(n^2),$$

$$(A.13) \quad \sum_{v \in S_{2n,2n}^G} d_v = 2 \sum_{j=n}^{3n/2-1} j + \sum_{j=0}^{n-1} j = \frac{7}{4}n^2 + O(n),$$

where the first summand is for nodes $(i, n/2)^2$ for $i = 1, \dots, n$, and the second one for nodes $(n/2, j)^1$ for $j = 1, \dots, 8$. Consequently,

$$(A.14) \quad \mu_m(S_{2n,2n}^G) = \sum_{v \in S_{2n,2n}^A} d_v(d_v + 3)/2 = \frac{23}{24}n^3 + O(n^2),$$

$$(A.15) \quad \mu_s(S_{2n,2n}^G) = \sum_{v \in S_{2n,2n}^A} (d_v + 1) = \frac{7}{4}n^2 + O(n).$$

We skip the detailed analysis but just list the results for A , B , and C below:

$$(A.16) \quad \sum_{v \in S_A} d_v = \sum_{j=3n/2}^{2n-1} j + \sum_{j=2n}^{5n/2-1} j + \sum_{j=2n}^{3n-1} j,$$

$$(A.17) \quad \mu_m(S_{2n,2n}^A) = \frac{125}{24}n^3 + O(n^2),$$

$$(A.18) \quad \mu_s(S_{2n,2n}^A) = \frac{9}{2}n^2 + O(n);$$

$$(A.19) \quad \sum_{v \in S_B} d_v = 2 \sum_{j=5n/2}^{3n-1} j + \sum_{j=3n}^{4n-1} j,$$

$$(A.20) \quad \mu_m(S_{2n,2n}^B) = \frac{239}{24}n^3 + O(n^2),$$

$$(A.21) \quad \mu_s(S_{2n,2n}^B) = \frac{25}{4}n^2 + O(n);$$

$$(A.22) \quad \sum_{v \in S_C} d_v = 2 \sum_{j=3n}^{7n/2-1} j + \sum_{j=4n}^{5n-1} j,$$

$$(A.23) \quad \mu_m(S_{2n,2n}^C) = \frac{371}{24}n^3 + o(n^3),$$

$$(A.24) \quad \mu_s(S_{2n,2n}^C) = \frac{31}{4}n^2 + o(n^2).$$

While $S_{2n,2n}^G$ and $S_{2n,2n}^C$ are eliminated in the order of the column nodes $(\cdot, n/2)^2$ followed by the row nodes $(n/2, \cdot)^1$, we eliminate $S_{2n,2n}^B$ in the order of the row nodes $(n/2, \cdot)^1$ followed by the column nodes $(\cdot, n/2)^2$. Moreover, we eliminate $S_{2n,2n}^A$ by

$$\begin{aligned} (i, n/2)^2 & \quad \text{for } i = 1, \dots, n/2, \\ (n/2, j)^1 & \quad \text{for } i = n/2 + 1, \dots, n, \\ (n/2, j)^1 & \quad \text{for } i = 1, \dots, n/2, \\ (i, n/2)^2 & \quad \text{for } i = n/2 + 1, \dots, n, \end{aligned}$$

in this order. This approach gives a better result.

Given $\mu(S_{2n,2n}^G)$, $\mu(S_{2n,2n}^A)$, $\mu(S_{2n,2n}^B)$, and $\mu(S_{2n,2n}^C)$, we can solve from the above recursions (see [18]) that the multiplication cost of applying CF to $G_{2n,2n}$ is $943n^3/84 + O(n^2 \log_2 n)$ and the storage cost is $(31/4)n^2 \log_2 n + O(n^2)$. Similar results can be derived for matrices whose adjacency graphs have the form shown in Figure 2.

Acknowledgments. We would like to thank Stanley Osher for introducing us to total variation image restoration and for many helpful and stimulating discussions, and Jinjun Xu for making his PDE codes available to us.

REFERENCES

- [1] F. ALIZADEH AND D. GOLDFARB, *Second-order cone programming*, Math. Program., 95 (2003), pp. 3–51.
- [2] F. ALIZADEH AND S. SCHMIETA, *Optimization with Semidefinite, Quadratic, and Linear Constraints*, Technical report, RUTCOR, Rutgers University, Piscataway, NJ, 1997.
- [3] K.D. ANDERSEN, E. CHRISTIANSEN, A.R. CONN, AND M.L. OVERTON, *An efficient primal-dual interior-point method for minimizing a sum of Euclidean norms*, SIAM J. Sci. Comput., 22 (2000), pp. 243–262.
- [4] G. AUBERT AND P. KORNPROBST, *Mathematical Problems in Image Processing*, Appl. Math. Sci. 147, Springer, New York, 2002.
- [5] A. CHAMBOLLE, *An algorithm for total variation minimization and applications*, J. Math. Imaging Vis., 20 (2004), pp. 89–97.
- [6] A. CHAMBOLLE AND P.-L. LIONS, *Image recovery via total variation minimization and related problems*, Numer. Math., 76 (1997), pp. 167–188.
- [7] R. CHAN, T. CHAN, AND H.-M. ZHOU, *Continuation Method for Total Variation Denoising Problems*, CAM Report 95-18, UCLA, 1995.
- [8] T. CHAN AND S. ESEDOGLU, *Aspects of Total Variation Regularized L^1 Function Approximation*, CAM Report 04-07, UCLA, 2004.
- [9] T.F. CHAN, G.H. GOLUB, AND P. MULET, *A nonlinear primal-dual method for total variation-based image restoration*, SIAM J. Sci. Comput., 20 (1999), pp. 1964–1977.
- [10] A. GEORGE, *Nested dissection of a regular finite element mesh*, SIAM J. Numer. Anal., 10 (1973), pp. 345–363.
- [11] D. GOLDFARB AND K. SCHEINBERG, *Numerically Stable Product-Form Cholesky (LDL^T) Factorization Based Implementations of Interior Point Methods for Second-Order Cone Programming*, CORC Report TR-2003-05, IEOR Department, Columbia University, New York, 2003.

- [12] M. HINTERMÜLLER AND K. KUNISCH, *Total bounded variation regularization as a bilaterally constrained optimization problem*, SIAM J. Appl. Math., 64 (2004), pp. 1311–1333.
- [13] M. HINTERMÜLLER AND G. STADLER, *An Infeasible Primal-Dual Algorithm for TV-Based inf-Convolution-type Image Restoration*, Technical report TR04-15, CAAM Department, Rice University, Houston, TX, 2004.
- [14] M. LYSAKER, A. LUNDERVOLD, AND X.-C. TAJ, *Noise removal using fourth-order partial differential equations with applications to medical magnetic resonance images in space and time*, IEEE Trans. Image Process., 12 (2003), pp. 1579–1590.
- [15] Y. MEYER, *Oscillating Patterns in Image Processing and Nonlinear Evolution Equations*, Univ. Lecture Ser. 22, AMS, Providence, RI, 2002.
- [16] S. OSHER, A. SOLÉ, AND L. VESE, *Image decomposition and restoration using total variation minimization and the H^{-1} norm*, Multiscale Model. Simul., 1 (2003), pp. 349–370.
- [17] I. POLLAK, A. WILLSKY, AND Y. HUANG, *Nonlinear evolution equations as fast and exact solvers of estimation problems*, IEEE Trans. Signal Process., 53 (2005), pp. 484–498.
- [18] D. ROSE AND G. WHITTEN, *A recursive analysis of dissection strategies*, in Sparse Matrix Computations, Academic Press, New York, 1976, pp. 59–84.
- [19] L. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, Phys. D, 60 (1992), pp. 259–268.
- [20] D. STRONG AND T. CHAN, *Edge-preserving and scale-dependent properties of total variation regularization*, Inverse Problems, 19 (2003), pp. 165–187.
- [21] L. VESE, *A study in the BV space of a denoising-deblurring variational problem*, Appl. Math. Optim., 44 (2001), pp. 131–161.
- [22] L. VESE AND S. OSHER, *Modeling textures with total variation minimization and oscillating patterns in image processing*, J. Sci. Comput., 19 (2003), pp. 553–572.
- [23] C.R. VOGEL AND M.E. OMAN, *Iterative methods for total variation denoising*, SIAM J. Sci. Comput., 17 (1996), pp. 227–238.
- [24] W. YIN, D. GOLDFARB, AND S. OSHER, *Total Variation-Based Image Cartoon-Texture Decomposition*, CORC Report, TR-2005-01, IEOR Department, Columbia University, New York, 2005.