SIAM J. SCI. COMPUT.
Vol. 30, No. 2, pp. 855–872

DEBLURRING METHODS USING ANTIREFLECTIVE BOUNDARY CONDITIONS*

MARTIN CHRISTIANSEN† AND MARTIN HANKE†

Abstract. In this paper we consider the numerical solution of self-adjoint deblurring problems on bounded intervals. For these problems it has recently been shown that appropriate modeling of the solution near the boundary of the interval may significantly improve the numerical reconstructions. Among the alternatives the so-called antireflective boundary condition appears to be the best known choice. Here we develop an appropriate, i.e., stable and efficient implementation of this model in two steps, namely by (i) transforming the problem to homogeneous boundary values, and (ii) using the fast sine transform to solve the transformed problem. This approach allows us to incorporate regularization in a very straightforward way. Numerical reconstructions are superior qualitatively and quantitatively to those obtained with the reblurring method of Donatelli and Serra-Capizzano.

Key words. deblurring problems, boundary conditions, fast sine transform, Tikhonov regularization

AMS subject classifications. 65F22, 65R32, 65T50

DOI. 10.1137/060671413

1. Introduction.

An important task in image reconstruction is the deblurring problem. In its simplest appearance this problem consists of approximating a function $f^* : \mathbb{R}^2 \to \mathbb{R}$–the original scene–from a blurred photo $g^* : \mathcal{I} \to \mathbb{R}$, where

$$g^*(x) = \int_{\mathbb{R}^2} k(x-y) f^*(y)\, dy\,, \qquad x \in \mathcal{I}\,. \tag{1.1}$$

Here, we consider $\mathcal{I} \subset \mathbb{R}^2$ to be a bounded rectangle and assume that the known characteristics of the imaging system are encoded in the so-called point spread function $k : \mathbb{R}^2 \to \mathbb{R}$; in particular, this implies that k is given. We will, moreover, assume throughout that k is *quadrantically symmetric*, i.e.,

$$k(x_1, x_2) = k(|x_1|, |x_2|)\,, \qquad x = (x_1, x_2) \in \mathbb{R}^2\,,$$

and has compact support containing the origin in its interior. Because of the compact support the integration involves only values of f^* over a bounded domain, but this domain may be considerably larger than $\mathcal{I}$, and hence problem (1.1) is underdetermined.

There are several possibilities for coping with this underdetermination; see, e.g., Hansen, Nagy, and O'Leary [6]. One way is to replace the integral in (1.1) by an integral over $\mathcal{I}$; this is equivalent to assuming that f^* is zero outside of $\mathcal{I}$ ("zero padding"). One can therefore imagine that this will yield reasonable reconstructions only when the true solution f^* has zero boundary values on $\partial\mathcal{I}$. Computationally, this approach leads to a linear system of equations (after discretizing the integral by a rectangular quadrature rule, say) which has a doubly Toeplitz structure. Many fast algorithms have been designed to deal with these highly structured but still nontrivial linear systems; cf. Ng [9].

*Received by the editors October 4, 2006; accepted for publication (in revised form) July 16, 2007; published electronically February 14, 2008.
http://www.siam.org/journals/sisc/30-2/67141.html
†Institut für Mathematik, Johannes Gutenberg-Universität Mainz, 55099 Mainz, Germany (martin-c-@web.de, hanke@math.uni-mainz.de).

Alternatively, when the boundary values of f^* are known to be the same on opposite sides of $\mathcal{I}$, one can base the reconstruction on the assumption that f^* is periodic with interval $\mathcal{I}$ of periodicity. While this assumption may be somewhat artificial in many circumstances, and will then cause so-called ringing artefacts, this leads to a linear system of equations with a doubly circulant structure that can be solved very efficiently with fast Fourier transform (FFT) techniques.

Yet another option is to model f^* by a function that is even with respect to the edges of $\mathcal{I}$; cf., e.g., Ng, Chan, and Tang [10]. This model has the advantage over the previous ones that a continuous scene $f^*|_{\mathcal{I}}$ always yields a continuous extension to all of $\mathbb{R}^2$. Therefore the approximation error is typically smaller than for zero padding—without any further requirements on the boundary values of f^*. The resulting discrete problems can be solved almost as cheaply as in the periodic case using fast cosine transforms. Note that the corresponding reconstructions of f^* are (discrete) functions with a vanishing normal derivative on $\partial\mathcal{I}$, which somehow limits the degree of approximation that can be expected from this approach.

In retrospect it is quite surprising that such a considerable amount of the literature deals with the zero padding model and the resulting doubly Toeplitz systems of equations. For, given zero boundary values of f^* on $\partial\mathcal{I}$, the true scene can also be approximated by a function f which is odd with respect to the edges of $\mathcal{I}$, and from the approximation theory point of view the modeling error will generically be even smaller than in the even case. If the point spread function is quadrantically symmetric, an efficient implementation of this approach can be based on the fast sine transform (FST). For these reasons we follow Serra-Capizzano [11] and strongly advocate this idea. We return to this model in sections 2 and 3.

While this last approach is restricted to homogeneous boundary values of f^*, there exists a modification for inhomogeneous boundary values, namely the so-called *antireflective boundary conditions* suggested in [11]. In this model f^* is extended at $x \in \partial\mathcal{I}$ in the normal direction by a function which is odd with respect to a fictitious origin in the point $(x, f^*(x))$. This yields about the same approximation error as before, and the resulting linear system is again amenable to solution methods based on the FST.

Since the deblurring problem is a first kind integral equation, its solution is highly susceptible to data errors (which are inevitable), and thus problem (1.1) requires some kind of regularization. It turns out that Tikhonov regularization, to pick out one possibility, can be incorporated in a straightforward way into the aforementioned algorithms without significant loss of performance—except for the antireflective boundary condition model; see section 4.

This loss of performance was first observed by Donatelli and Serra-Capizzano [4]. To fix the problem, they suggested an ad hoc modification of Tikhonov regularization, which lacks a theoretical foundation up to now: neither its stability nor its convergence are yet understood. In section 4 we therefore propose a different regularization scheme, which comes in two steps. In the first step we apply a simple transformation to the problem to achieve homogeneous boundary values on ∂I for both the true scene and its image. In the second step, we use FSTs to solve the transformed problem.

Our numerical results in section 6 will provide evidence that the corresponding numerical reconstructions of f^* are superior to the ones from [4], while the implementation is at the same time somewhat easier conceptually and has the same complexity. Moreover, a theoretical investigation of stability and convergence can follow the well-developed lines of linear regularization theory; see, e.g., [5].

For simplicity we restrict our attention mostly to 1D problems, where $\mathbb{R}$ instead of $\mathbb{R}^2$ is the domain of integration, and $\mathcal{I}$ is a bounded interval. The algorithm nonetheless extends quite naturally to 2D problems, and we will briefly summarize this in section 7. We also prefer to stick to the continuous setting (formulating the problem within the space of square integrable functions) when setting the stage in section 2, because we feel that in this context a number of notations and phrasings appear to be more natural than in the fully discrete formulation that we will turn to from section 3 onwards.

2. Modeling aspects. In one space dimension all the models considered in the introduction can be formulated within the following mathematical framework. Let $k \in \mathcal{L}^\infty(\mathbb{R})$ be an even, nonnegative point spread function with compact support, and consider the spatially invariant blurring operator

$$K : \mathcal{L}^2(\mathbb{R}) \to \mathcal{L}^2(\mathcal{I}), \qquad Kf(x) = \int_{\mathbb{R}} k(x-y)f(y)\,dy, \quad x \in \mathcal{I}, \tag{2.1}$$

where we fix $\mathcal{I} = [0, \pi]$ for convenience. We denote by $f^* \in \mathcal{L}^2(\mathbb{R})$ the true (unknown) scene, and assume that we are given a blurred and noisy image g of f^* such that

$$\|g - Kf^*\|_{\mathcal{L}^2(\mathcal{I})} \le \delta. \tag{2.2}$$

We denote by $\delta/\|Kf^*\|_{\mathcal{L}^2(\mathcal{I})}$ the (relative) *noise level* in the given data.

Next, we choose a closed subspace $X_\pi \subset \mathcal{L}^2(\mathbb{R})$ with the induced topology of $\mathcal{L}^2(\mathbb{R})$ as a *model*, out of which approximations $f_\pi \approx f^*$ are to be selected; we restrict ourselves to approximations that depend linearly on the given data, i.e.,

$$f_\pi = R_{\pi,\alpha} g, \tag{2.3}$$

where $R_{\pi,\alpha}$ depends on the choice of X_π and on a regularization parameter $\alpha > 0$. For example, if Tikhonov regularization is applied to the approximate identity

$$K_\pi f = g, \qquad K_\pi = K|_{X_\pi} : X_\pi \to \mathcal{L}^2(\mathcal{I}),$$

then

$$R_{\pi,\alpha} = K_\pi^*(K_\pi K_\pi^* + \alpha I)^{-1}.$$

It is obvious that in this case f_π will belong to the range of K_π^* and thus to X_π.

The four models considered in the introduction can be embedded into this setting as follows:

1. *zero padding*: $X_\pi = \{ f \in \mathcal{L}^2(\mathbb{R}) : f(x) = 0 \text{ a.e. in } \mathbb{R} \setminus \mathcal{I} \}$.
2. *periodic continuation*: $X_\pi = \{ f \in \mathcal{L}^2(\mathbb{R}) : f(x) = f(x-\pi) \text{ a.e.} \}$.
3. *even continuation*: $X_\pi = \{ f \in \mathcal{L}^2(\mathbb{R}) : f(-x) = f(x) = f(2\pi - x) \text{ a.e.} \}$.
4. *odd continuation*:
$$X_\pi = \{ f \in \mathcal{L}^2(\mathbb{R}) : f(-x) = -f(x) = f(2\pi - x) \text{ a.e.} \}. \tag{2.4}$$
5. *antireflective continuation*:
$$X_\pi = \left\{ \begin{array}{c} f \in \mathcal{L}^2(\mathbb{R}) : \text{ there is } c_0 = c_0(f) \text{ and } c_\pi = c_\pi(f) \text{ with} \\ f(-x) + f(x) = c_0 \text{ and } f(2\pi - x) + f(x) = c_\pi \text{ a.e.} \end{array} \right\}. \tag{2.5}$$
Note that if $f \in X_\pi$ is continuous, then necessarily $c_0(f) = 2f(0)$ and $c_\pi(f) = 2f(\pi)$.

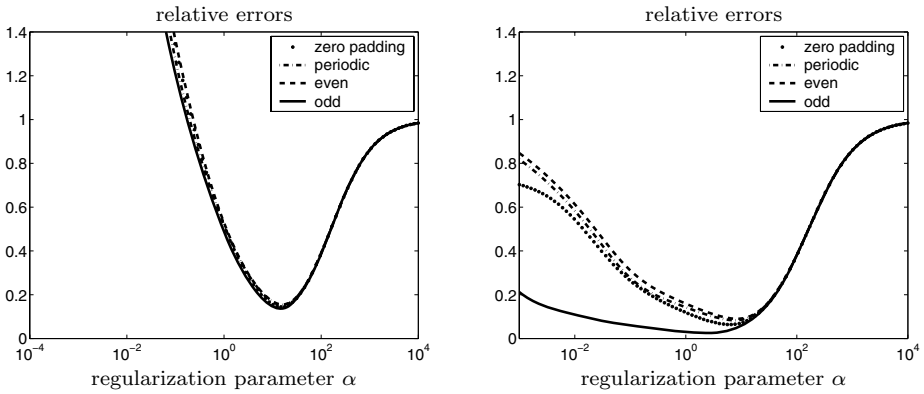

FIG. 2.1. *Regularization errors with* 10% *(left) and* 0.1% *(right) noise.*

For practical computations, X_π will actually be a finite-dimensional subspace of the respective sets, and this will therefore be preassumed in what follows.

If we introduce

$$f_\pi^* \;=\; \arg\min_{f_\pi \in X_\pi} \|f^* - f_\pi\|_{\mathcal{L}^2(\mathcal{I})} \qquad \text{and} \qquad g_\pi^* \;=\; K f_\pi^* \;=\; K_\pi f_\pi^*\,, \tag{2.6}$$

i.e., the best approximation $f_\pi^* \in X_\pi$ of the true scene $f^*|_{\mathcal{I}}$ and its image under K, then the following general error estimate for the approximation (2.3) is straightforward:

$$\begin{aligned} f_\pi - f^* \;&=\; f_\pi - f_\pi^* + f_\pi^* - f^* \\ &=\; R_{\pi,\alpha}(g - g_\pi^*) \;+\; (R_{\pi,\alpha} g_\pi^* - f_\pi^*) \;+\; (f_\pi^* - f^*)\,. \end{aligned} \tag{2.7}$$

The three terms in the second line of (2.7) correspond, from left to right, to a propagated data error, the approximation error of the regularization scheme, and the modeling error due to the choice of X_π. While the last term is independent of α, the first two terms counteract, and have to be balanced by a proper choice of the regularization parameter. The optimal balance depends on the rate of convergence of the regularization scheme, which in turn depends to some extent on some abstract "smoothness" of f_π^*; cf. [5].

We shall illustrate this for a specific example.

Example 2.1. Let f^* be the polynomial

$$f^*(x) \;=\; -7x^4 + 12x^3 - 6x^2 + x$$

with homogeneous boundary values on the interval $\mathcal{I} = [0, 1]$, and consider a motion blur g of f^* sampled on 129 equidistant pixel points within this interval. The motion blur involves 13 pixels; i.e., the kernel function k in (2.1) is the characteristic function of the interval $[-6/128, 6/128]$.

Figure 2.1 plots the relative errors (in the Euclidean norm) of Tikhonov regularization versus the regularization parameter for the different models from above; note that odd and antireflective continuations are the same for this example, as the true scene has homogeneous boundary data. Without regularization the resulting system matrices are mildly ill-conditioned, the condition number being somewhere between 500 and 1000.

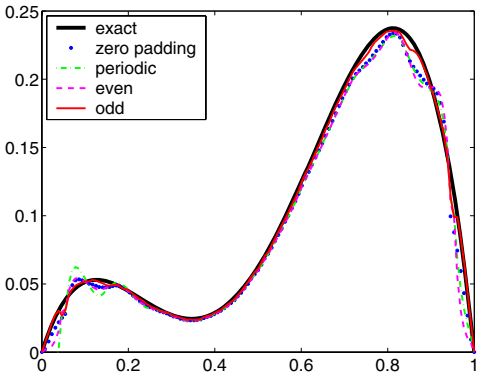

FIG. 2.2. *Reconstructions of f^* (0.1% noise).*

The two plots in Figure 2.1 correspond to two different noisy copies of Kf^*: In the left-hand plot the noise level $\|g - Kf^*\|_2/\|Kf^*\|_2$ is as large as 10%, whereas in the right-hand plot it is down to 0.1%. Our discussion above and these numerical results allow the following conclusions:

- When δ is large, the propagated data error dominates, and the modeling error is less important. Although the error with optimal regularization might slightly decrease with increasing smoothness of f_π^*, all models perform pretty much the same in our example.
  We mention that this is not quite the case for the reblurring scheme of [4], when the true scene fails to have zero boundary values. In fact, Donatelli and Serra-Capizzano report in [4] that the reblurring method deteriorates somewhat for large noise levels. Similar results were obtained in [3]; see also the subplot of Figure 6.3 below that corresponds to Example 6.2.
- When δ is small, however, the optimal error mainly depends on the choice of X_π, since $\|f_\pi^* - f^*\|_{\mathcal{L}^2(\mathcal{I})}$ is essentially a lower bound for the total error. In fact, in our example the right-hand plot in Figure 2.1 would not show any difference if no synthetic data error were superposed. In other words, the optimal error depends only on the distance between f^* and f_π^*, and between Kf^* and Kf_π^*, and thus only on the amenability of approximating f^* by functions from X_π. In our example this favors "odd continuation," because this is the only model which leads to C^1 approximations of the true polynomial f^*.

Figure 2.2 shows the reconstructions obtained by the four methods in the small noise case. Note that all but the one based on odd continuation show significant artefacts near the boundaries, being somewhat less pronounced for the zero padding approximation. Additionally, even in the interior of the interval, the "odd reconstruction" has fewer "wiggles" than the other three. This is so because the approximate data g_π^* are closer to Kf^* for this model, adding less "model noise" to the data.

3. Odd continuation. We now turn to a more detailed presentation of the numerical reconstruction algorithm for the odd continuation model. For this we shall assume that k is supported in an interval somewhat smaller than $[-\pi, \pi]$; this is not a restrictive assumption in practice as, for example in image reconstruction, the point-spread function (effectively) has a much smaller support than the image. Let X_π be chosen as in (2.4), i.e., every function $f_\pi \in X_\pi$ is odd (with respect to the origin) and

2π-periodic. Because of this we obtain for $x \in [0, \pi]$

$$\begin{aligned} Kf_\pi(x) &= \int_{\mathbb{R}} k(x-y) f_\pi(y)\, dy \\ &= \sum_{j\in\mathbb{Z}} \Big(\int_{2j\pi}^{(2j+1)\pi} k(x-y) f_\pi(y-2j\pi)\, dy \\ &\qquad\qquad - \int_{(2j-1)\pi}^{2j\pi} k(x-y) f_\pi(2j\pi - y)\, dy \Big) \\ &= \int_0^\pi \sum_{j\in\mathbb{Z}} \Big(k(x-y-2j\pi) - k(x+y-2j\pi) \Big) f_\pi(y)\, dy \\ &= \int_0^\pi \big(k(x-y) - k(x+y-2\pi) - k(x+y) \big) f_\pi(y)\, dy\,, \end{aligned}$$

where we have used the size restriction for the support of k in the last equality. Note that for every pair (x, y) at most two of the three terms of the kernel function will be nonzero.

The standard (although not necessarily best) discretization of these problems starts with an equidistant grid with grid size $h = \pi/(n+1)$, n some natural number, and defines the scaled values

$$k_j \;=\; h\, k(jh)\,, \qquad j \in \mathbb{Z}\,,$$

of the point spread function. Note that $k_j = 0$ for $|j| \geq n$ because of our restriction on the support of k. Evaluating Kf_π at $x = ih$, $i = 1, \ldots, n$, we obtain

$$Kf_\pi(ih) \;\approx\; \sum_{j=1}^{n} (k_{i-j} - k_{i+j-2n-2} - k_{i+j}) f_j\,, \qquad i = 1, \ldots, n\,,$$

where $f_j = f_\pi(jh)$; recall that $f_\pi(0) = f_\pi(\pi) = 0$. Forcing the right-hand side to be equal to $g_i = g(ih)$, we obtain the linear system

$$A\mathbf{f} = \mathbf{g}\,, \tag{3.1}$$

where $\mathbf{f} = [f_1, \ldots, f_n]^T$, $\mathbf{g} = [g_1, \ldots, g_n]^T$, and A is the difference of a Toeplitz and a Hankel matrix,

$$A = \begin{bmatrix} k_0 & k_1 & & \ldots & k_{n-2} & k_{n-1} \\ k_1 & k_0 & k_1 & & & k_{n-2} \\ \vdots & k_1 & k_0 & \ddots & & \vdots \\ \vdots & & \ddots & \ddots & \ddots & k_2 \\ k_{n-2} & & & \ddots & k_0 & k_1 \\ k_{n-1} & k_{n-2} & \ldots & k_2 & k_1 & k_0 \end{bmatrix} - \begin{bmatrix} k_2 & k_3 & \ldots & k_{n-1} & 0 & 0 \\ k_3 & & \iddots & 0 & 0 & 0 \\ \vdots & \iddots & \iddots & \iddots & 0 & k_{n-1} \\ k_{n-1} & 0 & \iddots & \iddots & \iddots & \vdots \\ 0 & 0 & 0 & \iddots & & k_3 \\ 0 & 0 & k_{n-1} & \ldots & k_3 & k_2 \end{bmatrix}. \tag{3.2}$$

Since k is even, we have $k_{-i} = k_i$ for all $i \in \mathbb{Z}$, and hence A is a symmetric matrix.

It has been shown by Bini and Capovani [1] and again by Boman and Koltracht [2] that matrices of the form (3.2) can be diagonalized as

$$A = SDS^T, \tag{3.3}$$

where D is a diagonal matrix containing the eigenvalues of A, and

$$S = \sqrt{\frac{2}{(n+1)}} \left[\sin \frac{jk\pi}{n+1} \right]_{j,k=1}^{n} \tag{3.4}$$

is the orthogonal sine matrix. Thus, the linear system (3.1) can be solved with only $O(n \log n)$ operations by using the FST; cf., e.g., Van Loan [12]. The same holds true if Tikhonov regularization is applied to problem (3.1).

For the reader's convenience we include another proof of this result by verifying that all columns $\mathbf{x}_k$ of S, $k = 1, \ldots, n$, are indeed eigenvectors of A. To this end, we rewrite (3.2) as

$$A = A_T - A_H$$

and introduce the antidiagonal unit matrix

$$J = \begin{bmatrix} & & 1 \\ & \iddots & \\ 1 & & \end{bmatrix} \in \mathbb{R}^{n \times n}.$$

Note that $J^2 = J^T J = I$, $JA_T J = A_T$, and $JA_H J = A_H$. Therefore, if we construct the block matrix

$$B = \begin{bmatrix} k_0 & \mathbf{b}^T & 0 & (J\mathbf{b})^T \\ \mathbf{b} & A_T & J\mathbf{b} & JA_H \\ 0 & (J\mathbf{b})^T & k_0 & \mathbf{b}^T \\ J\mathbf{b} & JA_H & \mathbf{b} & A_T \end{bmatrix} \qquad \text{with} \qquad \mathbf{b} = \begin{bmatrix} k_1 \\ \vdots \\ k_{n-1} \\ 0 \end{bmatrix}, \tag{3.5}$$

then we find that

$$B \begin{bmatrix} 0 \\ \mathbf{x} \\ 0 \\ -J\mathbf{x} \end{bmatrix} = \begin{bmatrix} 0 \\ A\mathbf{x} \\ 0 \\ -JA\mathbf{x} \end{bmatrix} \qquad \text{for every } \mathbf{x} \in \mathbb{R}^n. \tag{3.6}$$

In particular, this shows that if $\mathbf{x}$ is an eigenvector of A, then

$$\hat{\mathbf{x}} = \begin{bmatrix} 0 \\ \mathbf{x} \\ 0 \\ -J\mathbf{x} \end{bmatrix} \tag{3.7}$$

is an eigenvector of B for the same eigenvalue. By construction, B is a real symmetric $(2n+2) \times (2n+2)$ dimensional circulant matrix. Being circulant, the $(2n+2)$-dimensional Fourier vectors

$$\mathbf{z}_k = \left[e^{ijk\pi/(n+1)} \right]_{j=0}^{2n+1}, \qquad k = 0, \ldots, 2n+1,$$

are eigenvectors of B; cf., e.g., [6]. Moreover, as B is symmetric and real, the imaginary parts $\hat{\mathbf{x}}_k$ of $\mathbf{z}_k$ are also eigenvectors of B. Since $\hat{\mathbf{x}}_k$, $k = 1, \ldots, n$, are connected to the sine vectors $\mathbf{x}_k$ via (3.7), it follows from (3.6) that $\mathbf{x}_k$ is indeed an eigenvector of A.

Therefore every matrix A of (3.2) can be rewritten as in (3.3). Since it is easy to see that matrices of either form (3.2) or (3.3), respectively, form vector spaces of dimension n, we conclude that the two sets of matrices are the same.

4. Antireflective continuation. Given X_π of (2.5), i.e., the antireflective model, the resulting system $K_\pi f_\pi = g$ with $f_\pi \in X_\pi$ and $K_\pi = K|_{X_\pi}$ can be discretized in much the same way as in the previous paragraph, the only difference being that i and j now run from 0 to $n+1$, since the boundary values $f_0 = f(0)$ and $f_{n+1} = f(\pi)$ are also unknown. Accordingly, if we let $\mathbf{f}_{AR} = [f_0, \ldots, f_{n+1}]^T$ and $\mathbf{g}_{AR} = [g_0, \ldots, g_{n+1}]^T$, then this leads to the augmented system

$$A_{AR}\mathbf{f}_{AR} = \mathbf{g}_{AR}\,, \tag{4.1}$$

where the matrix

$$A_{AR} = \begin{bmatrix} s_0 & \mathbf{0} & 0 \\ \mathbf{s} & A & J\mathbf{s} \\ 0 & \mathbf{0} & s_0 \end{bmatrix} \tag{4.2}$$

contains the previous matrix A of (3.2) in its center block; the other nonzero entries of A_{AR} are given by

$$s_i = k_i + 2\sum_{j=i+1}^{n-1} k_j\,, \quad i = 0, \ldots, n-1, \qquad \text{and} \qquad \mathbf{s} = \begin{bmatrix} s_1 \\ \vdots \\ s_{n-1} \\ 0 \end{bmatrix}.$$

In particular, because of our assumptions on k,

$$s_0 = \sum_{j=1-n}^{n-1} k_j \tag{4.3}$$

is the total mass of the blur. We refer to [11] for a derivation of (4.2).

Linear systems of the form (4.1) can be solved in $O(n \log n)$ operations using the FST in a sophisticated way; cf. [11]. The same holds true for the Tikhonov regularized problem

$$(A_{AR}^T A_{AR} + \alpha I)\mathbf{f}_{AR} = A_{AR}^T \mathbf{g}_{AR}\,. \tag{4.4}$$

To see this, we observe that the central block of

$$A_{AR}^T A_{AR} + \alpha I = \begin{bmatrix} s_0^2 + \mathbf{s}^T\mathbf{s} + \alpha & \mathbf{s}^T A & \mathbf{s}^T J\mathbf{s} \\ A\mathbf{s} & A^2 + \alpha I & AJ\mathbf{s} \\ \mathbf{s}^T J\mathbf{s} & \mathbf{s}^T JA & s_0^2 + \mathbf{s}^T\mathbf{s} + \alpha \end{bmatrix}$$

can again be diagonalized with the sine matrix S. Thus, permuting this block to the (1,1)-position, the Tikhonov regularized problem can be solved efficiently using a Schur complement approach.

However, applying Tikhonov regularization to problem (4.1) turns out to be a pitfall similar to those discussed in a seminal paper by Varah [13] back in 1983: Solutions $\mathbf{f}_{AR} = [f_0, \ldots, f_{n+1}]^T$ of (4.4) necessarily belong to the range of A_{AR}^T, and accordingly, by virtue of (4.2), the inner components $f_1, \ldots, f_n$ of $\mathbf{f}_{AR}$ belong to the range of A, i.e., are given by a linear combination of discrete sine functions. As a consequence, these components call for a natural ("continuous") extension $f_0 = f_{n+1} = 0$, while on the other hand, generic elements from the range of A_{AR}^T will have nonzero boundary values.

We therefore expect to see severe boundary artefacts when Tikhonov regularization is applied to (4.1), and, in fact, this has been observed numerically by Donatelli and Serra-Capizzano [4]. Because of this, they recommend replacing the Tikhonov regularized system by the equation

$$(A_{AR}^2 + \alpha I)\mathbf{f}_{AR} = A_{AR}\mathbf{g}_{AR}\,, \tag{4.5}$$

i.e., substituting A_{AR} for A_{AR}^T in (4.4) and interpreting this as some kind of "reblurring" of the data. Again, (4.5) can be solved in $O(n \log n)$ operations by means of the FST. While the numerical results in [4] appear to support this approach, the theoretical properties remain dubious because the coefficient matrix in (4.5) is nonnormal (unless we face the trivial case where $k_i = 0$ for all $i \neq 0$). As of today neither a convergence nor a stability analysis appears to be within reach.

In the following section we shall therefore pursue a different idea to compute regularized solutions of (4.1). It will turn out that our algorithm has the same computational efficiency and may actually be somewhat simpler conceptually. The reconstructions appear to be at least as good as the ones from [4] (cf. section 6 for numerical examples), and a theoretical investigation of this approach on the basis of the known regularization theory (as, e.g., in [5]) is straightforward.

5. Regularization after transformation. From the particular form of A_{AR} it follows immediately that the spectrum of A_{AR} consists of the eigenvalues of A and twice the eigenvalue s_0. While this has already been observed in [11], we also need the eigenspace corresponding to $\lambda = s_0$ for our purposes below.

LEMMA 5.1. *The spectral radius of A_{AR} of* (2.5) *is given by the number s_0 of* (4.3). *The eigenspace corresponding to the eigenvalue $\lambda = s_0$ of the matrix A_{AR} contains the two vectors*

$$\mathbf{1} = \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix} \qquad \textit{and} \qquad \boldsymbol{\ell} = \begin{bmatrix} 0 \\ 1 \\ \vdots \\ n+1 \end{bmatrix}. \tag{5.1}$$

Proof. We will see in a minute that $\lambda = s_0$ is an eigenvalue of A_{AR}; hence the spectral radius of A_{AR} is at least as big as s_0. On the other hand, the spectral radius is bounded from above by the maximum absolute row sum norm of A_{AR}, and it is easy to see that this norm equals s_0. This shows that the spectral radius of A_{AR} is given by s_0.

Concerning the eigenvectors, we first observe that the discretization leading to (4.1) is such that the ith component g_i of $\mathbf{g}_{AR} = A_{AR}\mathbf{f}_{AR}$ satisfies

$$g_i \;=\; \sum_{j=1-n}^{n-1} k_j f_{i-j}\,, \qquad i = 0, \ldots, n+1\,,$$

provided that f_i, $i \in \mathbb{Z}$, are the entries of the antireflective continuation of $\mathbf{f}_{AR} = [f_0, \ldots, f_{n+1}]^T$.

For $\mathbf{f}_{AR} = \mathbf{1}$ this antireflective continuation yields $f_j = 1$ for all $j \in \mathbb{Z}$, and hence, $g_i = s_0$ for all $i = 0, \ldots, n+1$. In other words, we have $\mathbf{g}_{AR} = s_0\mathbf{1} = s_0\mathbf{f}_{AR}$; i.e., $\mathbf{1}$ is an eigenvector of A_{AR} corresponding to the eigenvalue $\lambda = s_0$.

For $\mathbf{f}_{AR} = \boldsymbol{\ell}$ the antireflective continuation yields $f_j = j$ for all $j \in \mathbb{Z}$, and hence,

$$g_i = \sum_{j=1-n}^{n-1} k_j f_{i-j} = \sum_{j=1-n}^{n-1} k_j (i-j) = i \sum_{j=1-n}^{n-1} k_j - \sum_{j=1-n}^{n-1} j k_j = i s_0 - \sum_{j=1-n}^{n-1} j k_j$$

for $i = 0, \ldots, n+1$. As $k_j = k_{-j}$ for all $j \in \mathbb{Z}$, the last term vanishes, which proves that $g_i = i s_0$, or again, that $\boldsymbol{\ell}$ is an eigenvector of A_{AR} corresponding to $\lambda = s_0$. ∎

We remark that there may be other linearly independent eigenvectors of A_{AR} for $\lambda = s_0$. In fact, the matrix

$$A = \begin{bmatrix} k_0 & -k_3 & 0 & k_3 & 0 \\ -k_3 & k_0 & 0 & 0 & k_3 \\ 0 & 0 & k_0 & 0 & 0 \\ k_3 & 0 & 0 & k_0 & -k_3 \\ 0 & k_3 & 0 & -k_3 & k_0 \end{bmatrix}$$

has another eigenvector $\mathbf{x}_4 = 1/2 \cdot [1, -1, 0, 1, -1]^T$ for $\lambda = s_0$, and hence the vector $[0, 1, -1, 0, 1, -1, 0]^T$ is an eigenvector of the associated matrix A_{AR} of (4.2) for the same eigenvalue.

Now we use Lemma 5.1 to introduce our transformation method. It has been observed in [11] that the boundary values of $\mathbf{f}_{AR}$ can readily be determined from the first and the last equations in (4.1), and this computation is stable: In fact, using the notation from the proof of Lemma 5.1, we have

$$f_0 = g_0 / s_0 \qquad \text{and} \qquad f_{n+1} = g_{n+1} / s_0 \,, \tag{5.2}$$

where the denominator s_0 is the total mass of the point spread function and thus well separated from zero.

Next, in contrast to [11], we use these boundary values of $\mathbf{f}_{AR}$ to transform the problem to a problem with homogeneous boundary conditions. We do this by subtracting an appropriate first degree polynomial. In the discrete setting, this amounts to the transformation

$$\mathbf{f}_H = \mathbf{f}_{AR} - f_0 \, \mathbf{1} - \frac{f_{n+1} - f_0}{n+1} \, \boldsymbol{\ell} \,. \tag{5.3}$$

In view of Lemma 5.1 the following transformation follows from (5.3) and (5.2):

$$\begin{aligned} \mathbf{g}_H = A_{AR} \mathbf{f}_H &= \mathbf{g}_{AR} - f_0 s_0 \, \mathbf{1} - \frac{f_{n+1} - f_0}{n+1} \, s_0 \, \boldsymbol{\ell} \\ &= \mathbf{g}_{AR} - g_0 \, \mathbf{1} - \frac{g_{n+1} - g_0}{n+1} \, \boldsymbol{\ell} \,. \end{aligned} \tag{5.4}$$

The first and last components of $\mathbf{f}_H$ and $\mathbf{g}_H$ are zero, and thus the first and last rows and columns of the transformed problem $A_{AR} \mathbf{f}_H = \mathbf{g}_H$ can be omitted, leaving us with the reduced problem

$$A \mathbf{f}'_H = \mathbf{g}'_H \tag{5.5}$$

for the inner components $\mathbf{f}'_H \in \mathbb{R}^n$ of $\mathbf{f}_H$, given the inner components $\mathbf{g}'_H \in \mathbb{R}^n$ of $\mathbf{g}_H$. Here, A is the matrix of (3.2), and the solution of (5.5) now might require some

kind of regularization, e.g., Tikhonov regularization. As mentioned in section 3, the regularized problem can be solved in $O(n \log n)$ operations using the FST.

Given an approximate solution $\tilde{\mathbf{f}}'_H \in \mathbb{R}^n$ of (5.5) and its extension $\tilde{\mathbf{f}}_H \in \mathbb{R}^{n+2}$ by zero boundary values, we finally obtain the desired approximation of $\mathbf{f}_{AR}$ from the back transformation

$$\tilde{\mathbf{f}}_{AR} = \tilde{\mathbf{f}}_H + f_0 \mathbf{1} + \frac{f_{n+1} - f_0}{n+1} \boldsymbol{\ell}, \tag{5.6}$$

where f_0 and f_{n+1} are given by (5.2).

6. Numerical results. We present some numerical results for two of the test problems from [4] to compare our transformation method with the reblurring method of [4]. In all plots, the broken line corresponds to the reblurring method, while the solid line corresponds to the new transformation method. Note that in both examples the true signal f^* (shown as thick black line) is only reconstructed in the subset of $(0, 1)$ where data have been sampled.

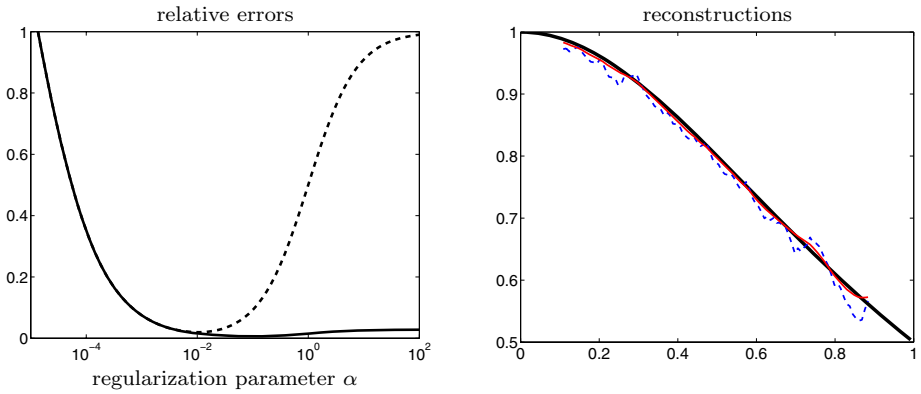

FIG. 6.1. *Example* 6.1 *(1% noise).*

Example 6.1. In the first problem the solution f^* is a smooth function,

$$f^*(x) = 1/(1 + x^2), \qquad x \in \mathbb{R},$$

and the kernel k is a truncated Gaussian; cf. [4, Table 3]. Observations of Kf^* are given on 101 equidistant points in the interval $[14/128, 114/128]$. The condition number of A is about $2.3 \cdot 10^4$.

Figure 6.1 compares the behavior of the two methods for a specific noise sample with $\|g - Kf^*\|_2 / \|Kf^*\|_2 = 0.01$, i.e., 1% noise: The left-hand plot shows the relative Euclidean errors of the reconstructions as a function of the regularization parameter α; the right-hand plot shows the optimal reconstructions of the two methods. In either case, the reconstruction has been found by optimizing the regularization parameter in such a way that the relative error of the reconstruction is minimal. The superiority of the transformation method is obvious, as its reconstruction is hard to distinguish from the true scene.

Example 6.2. The second example, taken from [4, Table 4], is well-posed (the condition number of A is below three), but the solution f^* is discontinuous. Being well-posed, we have run this problem with considerably more noise. Still, the reconstruction error is mainly influenced by the discontinuity of the solution, which causes Gibbs-like phenomena in both reconstructions; see Figure 6.2.

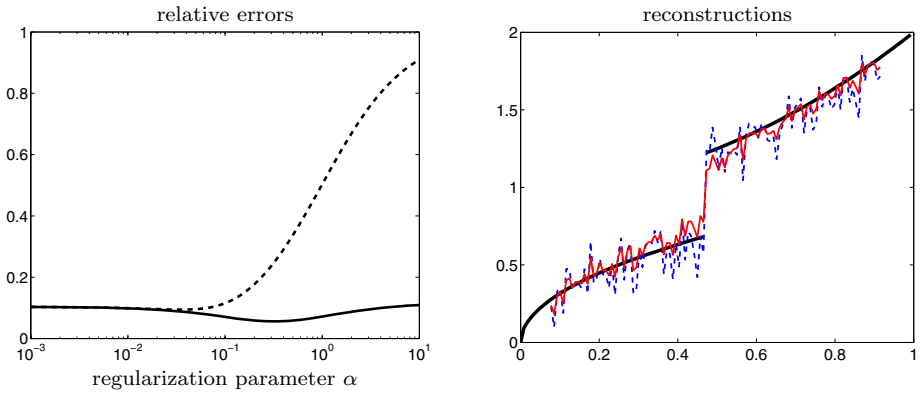

FIG. 6.2. *Example* 6.2 *(5% noise).*

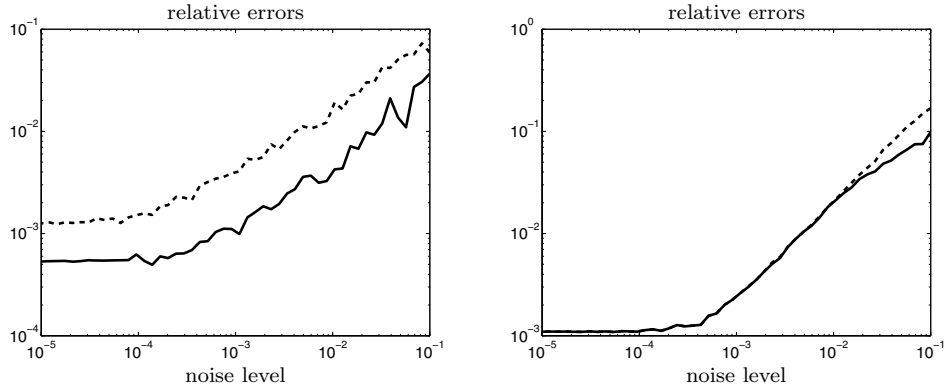

FIG. 6.3. *Examples* 6.1 *(left) and* 6.2 *(right).*

We now turn to a comparison of the two methods in dependence on the noise level. Figure 6.3 shows the optimal relative errors for the two examples and noise levels ranging from 10^{-5} up to 10%. The zigzags in the curves are due to the fact that the particular noise realizations are generated randomly; of course, for each noise level both methods use the same data to allow a fair comparison.

Recall that the reconstruction in Figure 6.2 obtained with the transformation method has fewer oscillations than that of the reblurring method. On the other hand, for smaller noise levels the two methods essentially yield approximations with the same quality, as can be deduced from the right-hand plot of Figure 6.3. This is different for the ill-posed example, though; see the left-hand plot in Figure 6.3. For this example, the reconstructions of the transformation method are consistently superior by a factor between two and five. As mentioned before, this is likely to be caused by the nonnormality of the coefficient matrix $A^2 + \alpha I$ in (4.5).

7. Extension to 2D problems. In image reconstruction (cf. (1.1)) f and g correspond to 2D images rather than 1D signals, and the discrete blurring coefficients

$$k_{ij} = h^2 k(ih, jh)\,, \qquad i, j \in \mathbb{Z},$$

fill up a 2D array. (Here, again, h denotes the mesh size of the grid.) Assuming antireflective boundary conditions on a bounded interval in the plane, the deconvolution problem leads to a linear system of equations for the unknown pixel values whose structure is again described in [11]. In what follows, we briefly comment on how the ideas from section 5 extend to this 2D situation.

As in the 1D case, the idea is to transform the given data to homogeneous boundary conditions. To this end, we gather the given pixel values in a matrix G_{AR}, corresponding to the size of the image. For notational simplicity we shall assume that this matrix is square, i.e., $G_{AR} \in \mathbb{R}^{(n+2)\times(n+2)}$. We can transform G_{AR} to homogeneous boundary values by setting

$$G_H \;=\; G_{AR} \;-\; \mathbf{1}\mathbf{a}_g^T \;-\; \boldsymbol{\ell}\mathbf{b}_g^T \;-\; \mathbf{c}_g\mathbf{1}^T \;-\; \mathbf{d}_g\boldsymbol{\ell}^T$$

with appropriate vectors $\mathbf{a}_g, \ldots, \mathbf{d}_g \in \mathbb{R}^{n+2}$ and with the vectors $\mathbf{1}$ and $\boldsymbol{\ell}$ from (5.1). (Note that each of the four correction terms is a dyadic vector product.) The homogeneous problem can be solved as in section 3 by using 2D FSTs, provided that k is quadrantically symmetric and that $k_{ij} = 0$ whenever $\max\{|i|, |j|\} \geq n$. We denote this solution by $F_H \in \mathbb{R}^{(n+2)\times(n+2)}$, using the same ordering of the pixel values as in G_H.

In a second step we need to backtransform F_H to obtain the approximation F_{AR} of the original image. Similar to (5.3), this backtransformation takes the form

$$F_{AR} \;=\; F_H \;+\; \mathbf{1}\mathbf{a}_f^T \;+\; \boldsymbol{\ell}\mathbf{b}_f^T \;+\; \mathbf{c}_f\mathbf{1}^T \;+\; \mathbf{d}_f\boldsymbol{\ell}^T$$

with $\mathbf{a}_f, \ldots, \mathbf{d}_f \in \mathbb{R}^{n+2}$, but in contrast to the 1D case the four correction terms no longer correspond to eigenvectors of the 2D antireflective coefficient matrix. Nonetheless, the vectors $\mathbf{a}_f$, $\mathbf{b}_f$, $\mathbf{c}_f$, and $\mathbf{d}_f$ can easily be computed, as they can be shown to solve the four linear systems

$$A_{AR}\mathbf{a}_f = \mathbf{a}_g\,, \quad A_{AR}\mathbf{b}_f = \mathbf{b}_g\,, \quad A_{AR}\mathbf{c}_f = \mathbf{c}_g\,, \quad A_{AR}\mathbf{d}_f = \mathbf{d}_g\,, \tag{7.1}$$

where $A_{AR} \in \mathbb{R}^{(n+2)\times(n+2)}$ is precisely the matrix (4.2) for the 1D problem with blurring coefficients

$$k_i = \sum_{j=1-n}^{n-1} k_{ij}\,, \qquad i = 1-n, \ldots, n-1\,. \tag{7.2}$$

We now give a brief sketch of the proof of the above transformation; the remaining details can be found in [3]. More precisely, we will show as an example that if $\mathbf{b}_f$ solves $A_{AR}\mathbf{b}_f = \mathbf{b}_g$, then the 2D antireflective blurring matrix maps the pixel vector $\boldsymbol{\ell}\otimes\mathbf{b}_f$ onto the pixel vector $\boldsymbol{\ell}\otimes\mathbf{b}_g$. Or, in an equivalent reformulation, if $A_{AR}\mathbf{b}_f = \mathbf{b}_g$, then the 2D space invariant blur with blurring coefficients k_{ij} maps the antireflective extension of the pixel matrix $\boldsymbol{\ell}\mathbf{b}_f^T$ onto $\boldsymbol{\ell}\mathbf{b}_g^T$. For the proof, let b_m, $m \in \mathbb{Z}$, be the entries of the antireflective extension of the vector $\mathbf{b}_f$; then the doubly infinite

antireflective extension $[b^f_{jm}]$ of the pixel matrix $\boldsymbol{\ell}\mathbf{b}_f^T$ has the form

$$\begin{array}{ccc|ccccc|cc}
 & & \vdots & \vdots & & & \vdots & \vdots & & \\
\cdots & & -b_{-1} & -b_0 & -b_1 & \cdots & -b_{n+1} & -b_{n+2} & \cdots \\
\hline
\cdots & & 0 & 0 & 0 & \cdots & 0 & 0 & \cdots \\
 & & b_{-1} & b_0 & b_1 & \cdots & b_{n+1} & b_{n+2} & \\
 & & \vdots & \vdots & & & \vdots & \vdots & \\
\cdots & & (n+1)\,b_{-1} & (n+1)\,b_0 & (n+1)\,b_1 & \cdots & (n+1)\,b_{n+1} & (n+1)\,b_{n+2} & \cdots \\
\hline
\cdots & & (n+2)\,b_{-1} & (n+2)\,b_0 & (n+2)\,b_1 & \cdots & (n+2)\,b_{n+1} & (n+2)\,b_{n+2} & \cdots \\
 & & \vdots & \vdots & & & \vdots & \vdots &
\end{array}$$

with the central part corresponding to the index range $0 \le j, m \le n+1$ being the original $(n+2) \times (n+2)$ matrix $\boldsymbol{\ell}\mathbf{b}_f^T$.

Applying now the space invariant blurring operator to this "image" yields a blurred image, or rather a doubly infinite pixel matrix $[b^g_{il}]$, whose entries are

$$\begin{aligned}
b^g_{il} &= \sum_{j,m=-\infty}^{\infty} k_{i-j,l-m} b^f_{jm} = \sum_{j,m=-\infty}^{\infty} k_{i-j,l-m} j b_m = \sum_{m=-\infty}^{\infty} \Big(\sum_{j=-\infty}^{\infty} j k_{i-j,l-m} \Big)\, b_m \\
&= \sum_{m=-\infty}^{\infty} \left(\sum_{j=-\infty}^{\infty} i k_{i-j,l-m} + \sum_{j=-\infty}^{\infty} (j-i)\, k_{i-j,l-m} \right) b_m\,, \qquad i,l \in \mathbb{Z}\,.
\end{aligned}$$

Note that all series are actually finite sums because only finitely many blurring coefficients are nonzero. Now the second sum within the brackets vanishes because $k_{i-j,l-m}$ is even with respect to its first index, so that

$$b^g_{il} = i \sum_{m=-\infty}^{\infty} \left(\sum_{j=-\infty}^{\infty} k_{i-j,l-m} \right) b_m = i \sum_{m=-\infty}^{\infty} k_{l-m} b_m\,, \qquad i,l \in \mathbb{Z}\,,$$

by virtue of (7.2) and the symmetry of the blur. For $0 \le l \le n+1$ the value of the series on the right-hand side is precisely the lth entry of $A_{AR}\mathbf{b}_f = \mathbf{b}_g$, and thus the inner $(n+2) \times (n+2)$ submatrix of the pixel matrix $[b^g_{il}]$ coincides with the dyadic product $\boldsymbol{\ell}\mathbf{b}_g^T$, as was to be shown.

We emphasize that in an actual implementation of this approach Tikhonov regularization should be utilized in both steps of the algorithm, i.e., for the solution of the 2D homogeneous problem and for the solution of the four 1D inhomogeneous problems (7.1). In our computations we have taken the same regularization parameter for each of these five subproblems, since this was most straightforward to implement and gave reasonable results.

To illustrate that our method is a viable alternative to image reconstruction methods based on the FFT or fast cosine transform, we use an image[1] with 256×256 pixels taken from the Berkeley Segmentation Dataset and Benchmark [7]. This image has a

[1] Training Image #253036: www.eecs.berkeley.edu/Research/Projects/CS/vision/bsds/BSDS300/html/dataset/images/gray/253036.html

TABLE 7.1
Example 7.1: *Relative errors of the reconstructions.*

Periodic	Even	Antireflective	Blurred data
0.1274	0.1246	0.0847	0.0861

TABLE 7.2
Example 7.2: *Relative errors of the reconstructions.*

Periodic	Even	Antireflective	Blurred data
0.0965	0.0474	0.0474	0.0738

relatively smooth background, which should be in favor of the higher approximation order models. For all reconstructions to follow, the regularization parameters have been optimized individually in order to minimize the relative errors in terms of the Euclidean, or rather, Frobenius norm. Although one can argue whether this is a reasonable criterion, it is definitely an objective one, and at the same time is the natural criterion for the particular Tikhonov penalty term that is associated with this kind of regularization. All displayed images use the same gray scale interval from 0 to 255; computed intensities outside this range have been cut off. We recommend turning to the electronic version of the paper for a comparison of the reconstructions, as the resolution on the screen is typically better.

Example 7.1. In our first example the image has been blurred by an averaging filter over 3×3 pixels, with the same weight $1/9$ attributed to all pixels. $1\,\%$ random noise has been added on top of the data; see Figure 7.1. Although this problem is only moderately ill-conditioned, the numerical results are interesting in that the antireflective reconstruction is the only one that is superior to the given data in terms of the relative Euclidean error measure; cf. Table 7.1. The reason for this is that the noise speckles are amplified in the other two reconstructions in order to sharpen the contours and the edges of the tree and the animals. In contrast, the compromise achieved with the antireflective model yields a better total error reduction. (Tuning the regularization parameter appropriately, the antireflective reconstruction can alternatively be brought to a good agreement with the even one.)

Example 7.2. Our second example is more ill-conditioned than the first one because this time the filter averages 11×11 pixels. On the other hand, the noise level is much smaller here ($0.05\,\%$); see Figure 7.2. All reconstructions show Gibbs effects along the contours of the tree and the animals; in addition, the periodic reconstruction exhibits strong ringing artefacts. This time the relative errors of the even and antireflective reconstructions happen to match almost exactly (cf. Table 7.2), and the two images agree pretty well, too. A more careful inspection of the even reconstruction reveals some minor Gibbs waves in the encircled region near the upper end of the reconstruction. These waves are less pronounced in the antireflective reconstruction because of the higher approximation order for this model.

original image

blurred and noisy image

periodic reconstruction

even reconstruction

antireflective reconstruction

FIG. 7.1. *Example* 7.1: 3×3 *blur with* 1% *noise.*

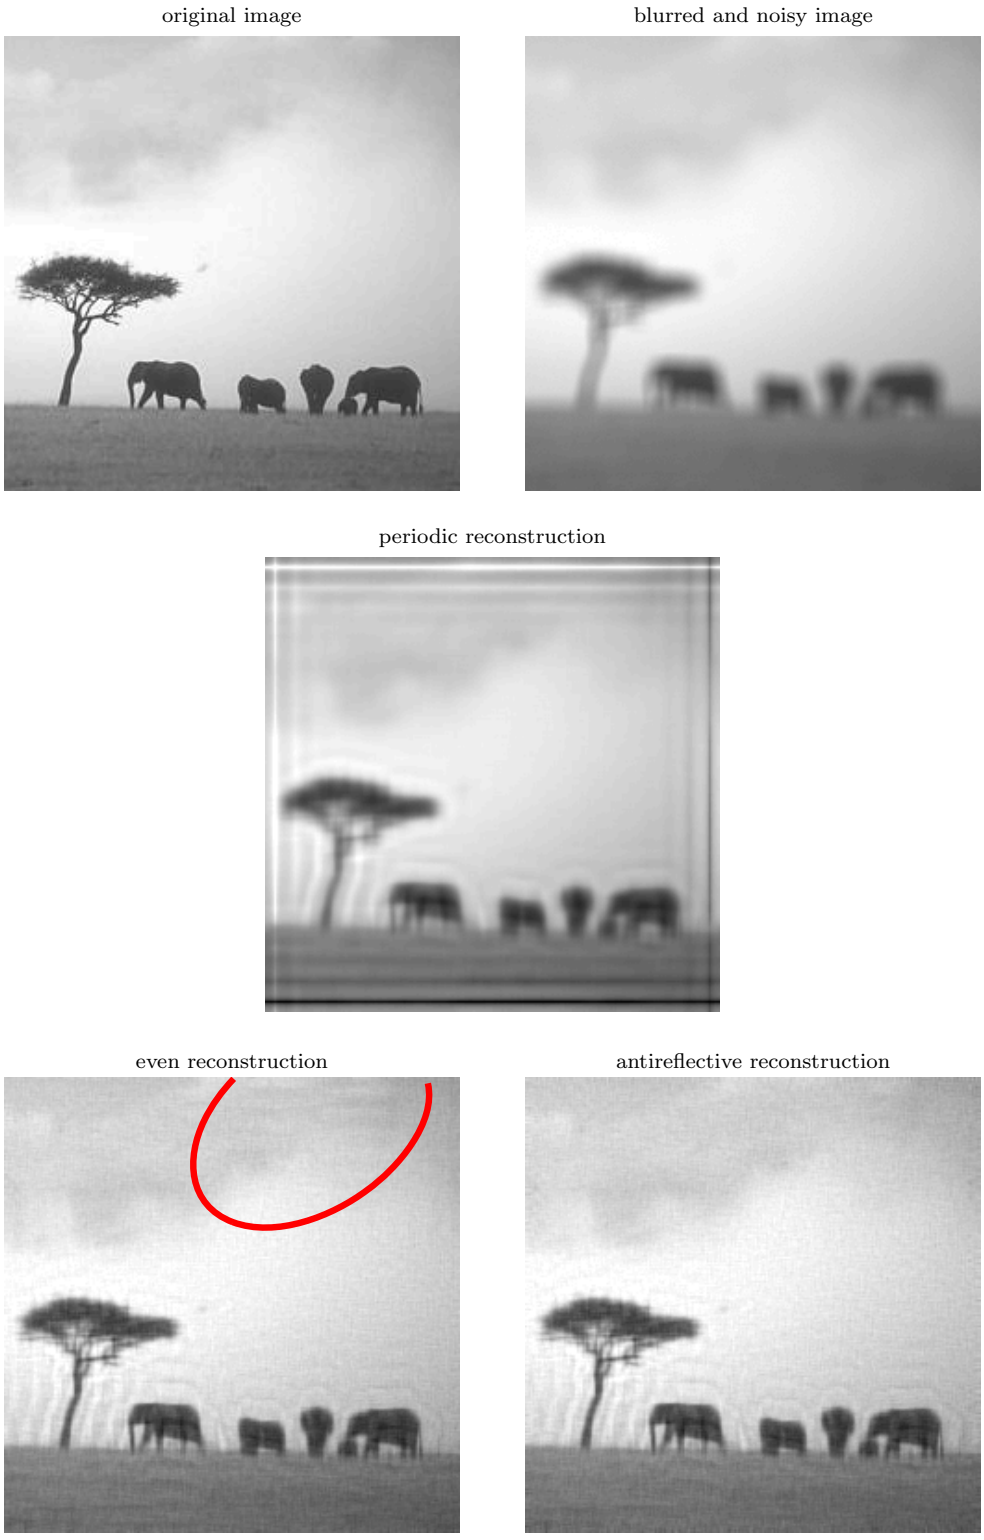

FIG. 7.2. *Example* 7.2: 11×11 *blur with* 0.05% *noise.*

8. Conclusion. We have investigated deblurring problems over $\mathbb{R}$ and $\mathbb{R}^2$, respectively, given only data on a bounded interval. Our arguments advocate the use of antireflective boundary conditions if the blur is quadrantically symmetric.

In case the given boundary data are homogeneous, the solution can be computed using fast sine transforms. For inhomogeneous boundary data we suggest transforming the problem into a homogeneous one, solving the homogeneous problem using fast sine transforms, and transforming back. The transformation involves only local computations with negligible computational overhead.

Regularization is easy to incorporate, which makes the method a vital alternative to other methods based on fast Fourier or cosine transforms.

Acknowledgments. We are indebted to Claus Schneider for pointing out to us that the regularized problem (4.4) can be solved in $O(n \log n)$ operations only.

Our implementation of the 2D image reconstruction codes greatly benefitted from the RestoreTools package[2] by Nagy, Palmer, and Perrone [8]. The authors also want to thank Jim Nagy for helpful advice on how to adapt this package to deal with antireflective boundary conditions.

Another acknowledgment goes to Per Christian Hansen for providing us with a pointer to the Berkeley Segmentation Dataset and Benchmark [7], and to Martin and Fowlkes for maintaining this dataset and making it available for noncommercial research and educational purposes.

REFERENCES

[1] D. Bini and M. Capovani, *Spectral and computational properties of band symmetric Toeplitz matrices*, Linear Algebra Appl., 52/53 (1983), pp. 99–126.

[2] E. Boman and I. Koltracht, *Fast transform based preconditioners for Toeplitz equations*, SIAM J. Matrix Anal. Appl., 16 (1995), pp. 628–645.

[3] M. Christiansen, *Die Behandlung antireflektiver Randbedingungen bei der numerischen Lösung von Faltungsgleichungen,* Diploma thesis, Institut für Mathematik, Johannes Gutenberg-Universität Mainz, Germany, 2006 (in German).

[4] M. Donatelli and S. Serra-Capizzano, *Anti-reflective boundary conditions and re-blurring*, Inverse Problems, 21 (2005), pp. 169–182.

[5] H. W. Engl, M. Hanke, and A. Neubauer, *Regularization of Inverse Problems*, Kluwer Academic Publishers, Dordrecht, The Netherlands, 1996.

[6] P. C. Hansen, J. G. Nagy, and D. P. O'Leary, *Deblurring Images: Matrices, Spectra, and Filtering*, Fund. Algorithms 3, SIAM, Philadelphia, 2006.

[7] D. Martin, C. Fowlkes, D. Tal, and J. Malik, *A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics*, in Proceedings of the 8th International Conference on Computer Vision, IEEE Press, Piscataway, NJ, 2001, Volume 2, pp. 416–423.

[8] J. G. Nagy, K. M. Palmer, and L. Perrone, *Iterative methods for image deblurring: A Matlab object oriented approach*, Numer. Algorithms, 36 (2004), pp. 73–93.

[9] M. K. Ng, *Iterative Methods for Toeplitz Systems*, Oxford University Press, Oxford, UK, 2004.

[10] M. K. Ng, R. H. Chan, and W.-C. Tang, *A fast algorithm for deblurring models with Neumann boundary conditions*, SIAM J. Sci. Comput., 21 (1999), pp. 851–866.

[11] S. Serra-Capizzano, *A note on antireflective boundary conditions and fast deblurring models*, SIAM J. Sci. Comput., 25 (2003), pp. 1307–1325.

[12] C. Van Loan, *Computational Frameworks for the Fast Fourier Transform*, SIAM, Philadelphia, 1992.

[13] J. M. Varah, *Pitfalls in the numerical solution of linear ill-posed problems*, SIAM J. Sci. Stat. Comput., 4 (1983), pp. 164–176.

[2] http://www.mathcs.emory.edu/∼nagy/RestoreTools