

ALGORITHMS AND SOFTWARE FOR ATMOSPHERIC IMAGE RECONSTRUCTION

W. K. Cochran, R. J. Plemmons and T. C. Torgersen

Wake Forest University, Box 7388, Winston-Salem, NC 27109

Abstract

A phase-diversity-based approach is taken in this paper for developing numerical techniques and a software package for postprocessing images taken through the atmosphere. The approach is based on a multiframe/multichannel formulation, where noisy, differently blurred short exposure images of the same object are available. Such situations occur when the same object is observed at successive time instants through the turbulent atmosphere, with a different transfer function at each instant (this can include the effect of partial compensation by adaptive optics). In general, the transfer functions differ primarily by changes in their phase. Phase-diversity image data is used to simultaneously estimate the object (i.e., the true image) and the phase, or wavefront profile. In our approach, regularization is applied and a resulting large-scale nonlinear optimization problem is solved efficiently using a limited memory quasi-Newton method with bound constraints. Some preliminary test results by applying our software package to simulated atmospheric image data are presented, and parallelization issues are addressed.

Keywords: space object imaging, image reconstruction, phase-diversity, regularization, optimization.

1 Introduction

Phase-diversity-based speckle (PDS) phase recovery and image reconstruction is an advanced imaging technique for restoring fine-resolution detail when imaging in the presence of phase aberrations such as atmospheric turbulence. The method, which in its basic form was developed by Gonsalves [3], has been applied using various numerical optimization techniques (see, e.g., [6, 11, 13, 15]), and its overall effectiveness for space-object identification has been well-accepted [11, 13]. The primary drawback to PDS in comparison the other image reconstruction schemes such as direct hardware implemented adaptive optics methods and deconvolution by wavefront sensing, has been the need for high performance optimization algorithm design and software implementation. Our purpose in this paper is to address these two important issues.

The PDS method involves the simultaneous collection of several short-exposure images, some of which are conventional images that have been blurred by unknown aberrations. One or more additional images are collected in *separate channels*, by blurring the first image by a known amount, e.g., using a beam splitter and an out-of-focus lens, which generates a quadratic blur. Using the images collected by the phase-diversity method as data, one can set up a mathematical optimization problem (typically using a maximum likelihood formulation) for recovering the original image as well as the phase aberrations.

Light rays propagating through the earth's atmosphere are distorted because of variations in the index of refraction due to differences in air temperature, humidity, and other factors. This causes distortion, or blurring. A consequence of this blurring is the limited spatial resolution of space objects viewed through ground-based telescopes. Two approaches are often taken to remove these degradations. With adaptive optics [8, 14], a deformable mirror is used to restore, or phase conjugate, the distorted light rays to planarity in real time, prior to the formation of the image of the object. A second approach is to solve the inverse problem of estimating the object and certain features of the atmosphere, given observed image data and an image formation model.

Modeling process of image formation is central to image reconstruction. In optical imaging it is generally assumed that quasi-monochromatic incoherent light energy propagates from object plane sources through an intervening medium and the telescope optical system, to an image plane. Let the object and image plane

irradiance distributions be denoted by $f(\nu, \mu)$ and $d(x, y)$, respectively. A simple model of the spatially-invariant image formation process (see, e.g., [9]) is the convolution integral equation

$$d(x, y) = \iint s(x - \nu, y - \mu) f(\nu, \mu) d\nu d\mu \odot \eta(x, y) = (s \star f)(x, y) \odot \eta(x, y), \quad (1)$$

where \odot indicates a point-by-point operation (e.g., addition or multiplication) with the noise process $\eta(x, y)$, and \star represents convolution. Here d thus represents *image data*, f represents the *object*, or true image, and s is known as the *point spread function*, or PSF. The PSF is the image that would result from an idealized point object, and it characterizes atmospheric blurring effects. In this context, the inverse problem is to determine both the PSF s and the object f , given the image data d . This inverse problem, which is referred to as *blind deconvolution*, is ill-posed in several respects. First, since convolution is symmetric (i.e., $s \star f = f \star s$), one cannot uniquely determine both s and f from a “single” observation of d . Additional information is needed. Even if the nonuniqueness difficulty in the model (1) can be overcome, the additional difficulty of *instability* may arise, since we are solving an ill-posed inverse problem [2]. The term *regularization* is used for schemes which restore stability in a manner which provides a good approximation to the true solution.

The atmospheric imaging problem we consider here has the form (1) with important additional special structure. The PSF can be described in terms of a function known as the *phase*, or wavefront profile. A physical process known as *phase-diversity* [3, 6, 11, 13, 15] can be applied to generate additional data to partially overcome the nonuniqueness difficulties in blind deconvolution. From several phase-diversity images, one then seeks to estimate both the object f and the phase.

The paper is outlined as follows. Our approach to PDS computations is described, regularization is discussed, and a derivation of the resulting cost functional is sketched in Section 2. A formula for the gradient of the cost functional is also given in this section, following the analysis in [15]. A fast limited memory quasi-Newton optimization method allowing bound constraints [1] for the minimization of our particular cost functional is the topic of Section 3. The method is a nonlinear minimization technique which combines low cost with rapid convergence. Only the cost functional and its gradient are computed at each iteration, allowing an efficient approximation to the Hessian based on updating. Our Fortran software package for multiframe PDS computations is discussed in Section 4, and parallel implementation issues are discussed in Section 5. In Section 6, we present some preliminary computational results obtained in applying the package to simulated atmospheric space object image data based on representative phase screens for distributed atmospheres. Finally, some future directions of this work are listed in Section 7.

2 Multiframe Phase-Diversity

For simplicity we first describe the single frame phase-diversity formulation, based on a “pair” of images. We will assume that recorded image data can be accurately represented as an $n_x \times n_y$ arrays of pixel intensities with components

$$[d]_{ij} = (s \star f)(x_i, y_j) \odot \eta_{ij}, \quad 1 \leq i \leq n_x, \quad 1 \leq j \leq n_y, \quad (2)$$

where where \odot indicates a point-by-point operation with the noise process $\eta(x, y)$, \star represents convolution, f denotes the true image, or object, and s denotes the point spread function, or PSF. The PSF s quantifies the blurring effects of the atmosphere. We will assume dependence on the phase, or wavefront profile, $\phi(x, y)$ as given by

$$s[\phi] = |\mathcal{F}^{-1}\{pe^{i\phi}\}|^2, \quad (3)$$

where \mathcal{F} denotes the 2-D Fourier transform, $i = \sqrt{-1}$, and $p = p(x, y)$ denotes the *pupil*, or aperture, function. For ground-based telescopes, p is an indicator function whose support is determined by the extent of the telescope mirror (see [9] for a detailed discussion of atmospheric imaging models).

The nonuniqueness and other difficulties for the model (2)-(3) can be at least partially resolved with phase-diversity data. In its simplest form, one collects a pair of images

$$\begin{aligned} d_1 &= s[\phi] \star f \odot \eta_1, \\ d_2 &= s[\phi + \theta] \star f \odot \eta_2, \end{aligned} \quad (4)$$

where θ represents a known phase perturbation (see [6]). In practice, this data is generated in hardware by splitting the beam of light collected from the telescope's primary mirror. From one beam, the conventional image d_1 is formed. An out-of-focus image d_2 is formed from the second beam. This corresponds to a quadratic phase perturbation, $\theta(x, y) = c(x^2 + y^2)$, where the constant c depends on the defocus length, generally chosen to have magnitude between π and 2π for space object imaging (see, e.g., [11]).

However, the single frame pair phase-diversity-based image recovery formulation described above will generally not lead to acceptable phase recoveries and image reconstructions [6, 9, 11]. The blind deconvolution problem (1) is still not adequately constrained. Fortunately, using multiple differently blurred frames of data is in itself a powerful constraint on the restored object.

PDS involves the collection of multiple short-exposure images of data, giving rise to a time series of recorded images. During short-exposure image acquisition, the turbulence structure of the atmosphere is effectively frozen, and the image is distorted by this instantaneous turbulence structure. These images are referred to as speckle because of their modulated appearance. Here, in spite of the severity of the image degradation, there is spatial-frequency information about the object up to the diffraction limit of the telescope system [9]. The high quality of PDS image reconstructions is illustrated in our tests reported in Section 5.

The multiframe/multichannel phase-diversity formulation can be described as follows. Let T denote the *number of time frames* of multiple short-exposure sets of data and K denote the *number of diversity channels for each time frame*. In applications relating to space object image recovery, T is generally a multiple of 16, since each Gemini record of data contains 16 frames (see [10]). Also, the number of diversity channels K for each time frame is usually taken to be 2 or 3, e.g., [11]. The the phase-diversity data is then described by the equation

$$d_{tk} = s[\phi_t + \theta_k] \star f \odot \eta_{tk}, \quad t = 1, \dots, T, \quad k = 1, \dots, K, \quad (5)$$

where $k = 1$ involves no diversity, i.e., θ_1 is a zero array for each ϕ_t .

Assuming a Gaussian fit-to-data criterion is applied, then the PDS blind deconvolution problem is to compute the phase screens ϕ_1, \dots, ϕ_T and the image f to minimize

$$J[\phi, f] = \frac{1}{2} \sum_{t=1}^T \sum_{k=1}^K \|s_{tk} \star f - d_{tk}\|^2 + \frac{\gamma}{2} J_{reg}[f] + \frac{\alpha}{2} J_{reg}[\phi_1, \dots, \phi_T], \quad (6)$$

where $s_{tk} = s[\phi_t + \theta_k]$.

In (6), $\alpha J_{reg}[\phi_1, \dots, \phi_T]$ is a regularization functional, whose purpose is to establish stability with respect to perturbations in the phase screens ϕ_t . Similarly, the term $\gamma J_{reg}[f]$ establishes stability with respect to perturbations in f . Here, γ and α are positive regularization parameters. Regularization is discussed in more detail in the next section.

2.1 Regularization

The PDS approach itself can greatly reduce the nonuniqueness difficulties. But it should be noted that constant offsets in ϕ still cannot be resolved, since $s[\phi + c] = s[\phi]$ for any fixed c . Such ambiguities include phase wrapping. This is a consequence of the fact that $e^{i\phi} = e^{i(\phi+2\pi)}$. One must also deal with instability with respect to perturbations in the data. This can be overcome by Tikhonov regularization [2], or penalized least squares. Regularization methods yield solutions to ill-posed inverse problems which depend on regularization parameters, which quantify the tradeoff between error amplification due to instability and truncation due to regularization. Regularization functionals can be used to enforce smoothness constraints or prior information about the unknowns [2, 4].

Given the multiframe/multichannel data (5), one might minimize the *joint* cost functional (6), where γ and α are nonnegative scalar regularization parameters, to be multiplied times regularization, or penalty, functionals J_{reg} . These penalty functionals restore stability in a manner which incorporates a priori information about the object and phase. For the purpose of simplifying the PDS computations, the object regularization functional in PDS methods is generally taken to be

$$J_{reg}[f] = \|f\|^2. \quad (7)$$

This corresponds to the minimal assumption that the object has finite intensity. It incorporates no prior smoothness assumptions.

There are various choices for the phase regularization functional. A study of three such choices is given in [4]. Define

$$J_{reg}[\phi_1, \dots, \phi_T] = \sum_{t=1}^T \|L\phi_t\|^2. \quad (8)$$

Then the choice of the operator L determines regularization functional for the phase. The choice $L = I$ corresponds to the minimal assumption that the phase has finite intensity. A smoothness constraint on the phase can be imposed by choosing L to be a differential operator, typically a Laplacian. However, if prior second order statistical information about the atmospheric turbulence is available, say from a wavefront sensor, then a sometimes more effective choice of L can be made [4]. If the phase ϕ is assumed to be a realization from a wide-sense stationary stochastic process whose covariance operator has a Von Karman spectrum [9] with covariance matrix C , then one might set

$$L = C^{-\frac{1}{2}}. \quad (9)$$

Choosing L as in (9) is called *MAP regularization*. Efficient numerical methods for computing C and the resulting L for this MAP regularization are given in [7, 8], where fast Hankel transforms are used. In ideal situations, the use of (9) for phase retrieval is recommended in the study by Irwan and Lane [4]. However, in working with real space object data, second order statistics may not be available, in which case the choice of L as the Laplacian may be most appropriate. We use this choice for L in our numerical tests reported in Section 6.

2.2 The Reduced Cost Functional and Its Gradient

Some of the notation and technical details in this section are taken from the paper by Vogel, Chan and Plemmons[15], which considered only the single frame phase-diversity case. Let upper case letters denote Fourier transformed variables. Then the multiframe/multichannel cost functional (6) has a Fourier domain representation

$$J_{\gamma, \alpha}[\phi, F] = \frac{1}{2} \left(\sum_{t=1}^T \sum_{k=1}^K \|S[\phi_t + \theta_k]F - D_{tk}\|^2 \right) + \frac{\gamma}{2} \|F\|^2 + \frac{\alpha}{2} J_{phase}[\phi_1, \dots, \phi_T], \quad (10)$$

where $J_{phase}[\phi_1, \dots, \phi_T]$ denotes the Fourier transform of the right-hand-side of (8). In order to simplify the notation, we set $S_{tk} = S[\phi_t + \theta_k]$ in the equations to follow, and in general, the subscript notation tk corresponds to a function evaluated at $\phi_t + \theta_k$. The exception is d_{tk} defined in (5), and its 2-D Fourier transform $D_{tk} = \mathcal{F}(d_{tk})$.

From (10) we thus obtain an unconstrained minimization problem with a very large number of unknowns. With $n_x \times n_y$ pixel image arrays, we have $(T + 1) n_x n_y$ unknowns. As in [3, 6, 15], we eliminate the object and cut the number of unknowns by $n_x n_y$. Setting $\frac{\partial J}{\partial F} = 0$ yields

$$F[\phi_1, \dots, \phi_T] = \frac{\sum_{t=1}^T \sum_{k=1}^K S_{tk}^* D_{tk}}{\gamma + \sum_{t=1}^T \sum_{k=1}^K |S_{tk}^*|^2}. \quad (11)$$

Here the superscript $*$ denotes complex conjugate, and $|\cdot|$ denotes component-wise magnitude of a complex quantity array. Given estimates for the T phase arrays ϕ_1, \dots, ϕ_T , one can take the inverse Fourier transform in (11) to obtain an estimate for the object f . Note that the positive object regularization parameter γ in the denominator of (11) induces stability by preventing division by very small quantities or zero.

By substituting $F = F[\phi_1, \dots, \phi_T]$ from (11) back into (10), one obtains the *reduced cost functional*

$$J[\phi_1, \dots, \phi_T] = \frac{1}{2} \left(\sum_{t=1}^T \sum_{k=1}^K \|D_{tk}\|^2 - \left\| \frac{\sum_{t=1}^T \sum_{k=1}^K S_{tk}^* D_{tk}}{\sqrt{\gamma + \sum_{t=1}^T \sum_{k=1}^K |S_{tk}^*|^2}} \right\|^2 \right) + \alpha J_{phase}[\phi_1, \dots, \phi_T]. \quad (12)$$

The gradient of the reduced cost functional then has a representation

$$g[\phi_1, \dots, \phi_T] = -2 \sum_{t=1}^T \sum_{k=1}^K \text{Imag}(H_{tk}^* \mathcal{F}(\text{Real}(h_{tk} \mathcal{F}^{-1}(V_{tk})))) + \alpha g_{reg}[\phi_1, \dots, \phi_T], \quad (13)$$

where $\theta_1 = 0$,

$$H_{tk} = p e^{i(\phi_t + \theta_k)}, \quad h_{tk} = \mathcal{F}^{-1}(H_{tk}), \quad (14)$$

$$s_{tk} = |h_{tk}|^2, \quad S_{tk} = \mathcal{F}(s_{tk}), \quad (15)$$

$$V_{tk} = F^* D_{tk} - |F|^2 S_{tk}, \quad (16)$$

and $g_{reg}[\phi_1, \dots, \phi_T]$ is the gradient of the phase regularization functional $J_{phase}[\phi]$. The latter quantity simplifies according to the choice of L in (8).

3 Limited Memory Constrained Optimization

To minimize the reduced cost functional $J[\phi_1, \dots, \phi_T]$ given in (12), we apply a limited memory quasi-Newton optimization method allowing bound constraints on the phase. The basic algorithm is given in [1] and has been incorporated into an optimization package at the Argonne National Laboratory in a highly efficient code.

Quasi-Newton methods yield approximations to a (local) minimizer u_* to the phase screens ϕ_t in (12) of the form

$$u_{i+1} = u_i + s_i, \quad i = 0, 1, \dots,$$

where $s_i = \mu d_i$ updates the current vector iterate, μ is a positive step length parameter, and the quasi-Newton direction vector d_i solves

$$H_i d = -g[u_i], \quad (17)$$

with H_i a symmetric positive definite approximation to the true Hessian $H[u_i]$, and $g[u_i]$ is the gradient vector given in (13), evaluated at the current phase approximation. Positive definiteness guarantees that d_i is a descent direction for J for some $\mu > 0$, provided the gradient is nonzero (see, e.g., [5]).

If H_i in (17) is taken to be the true Hessian H , then one obtains Newton's method. This method has the advantage of quadratic convergence near a local minimizer. Unfortunately, the computation, storage, and inversion of the Hessian may be prohibitively expensive. Moreover, far from a local minimizer the Hessian need not be positive definite, and hence the Newton step need not be a descent direction.

An alternative to Newton's method is the BFGS method (see, e.g., [1]). Given an initial Hessian approximation H_0 , it generates a sequence of Hessian approximates via the rank-two update

$$H_{i+1} = H_i + \frac{1}{y^T s} y y^T - \frac{1}{s^T H_i s} (H_i s)(H_i s)^T, \quad (18)$$

where $s = s_i = u_{i+1} - u_i$ is the current scaled step vector and $y = g[u_{i+1}] - g[u_i]$ is the difference between gradients. If H_i is positive definite and $y^T s$ is positive, then H_{i+1} is guaranteed to be positive definite. Under standard assumptions, i.e., J is smooth, u_0 is sufficiently close to a local minimizer u_* , H_0 is sufficiently close to $H[u_*]$, and $H[u_*]$ is positive definite, the BFGS method is rapidly convergent (see [1] for details).

A recursive formula for the inverses of the matrices in (18) is

$$H_{i+1}^{-1} = (I - \frac{1}{y^T s} s y^T) H_i^{-1} (I - \frac{1}{y^T s} y s^T) + \frac{1}{y^T s} s s^T. \quad (19)$$

This recursion can be used to easily solve equation (17) with a BFGS Hessian approximation H_i . To do so requires storage of i vectors s and y (i.e., all the previous steps and gradient differences), inner product computations involving the gradient $g[u_i]$ and the s and y vectors, and the computation of $H_0^{-1} v$ for some vector v . With the limited memory BFGS method, only a fixed number of the s and y vectors, say 5 pairs, are retained. As new vectors are added to storage, the oldest vectors are discarded. This can substantially reduce the storage requirements of the method and the cost of computing the quasi-Newton steps d .

The numerical optimization routine incorporated into our PDS software package is called L-BFGS-B, for limited memory BFGS with bound constraints. The basic code, in FORTRAN, is available, as is the paper [1], from the Argonne National Laboratory Optimization Technology Center on their web page at www-unix.mcs.anl.gov/neos/Server/solvers/BCO:L-BFGS-B/.

4 The Software Package

We have developed and tested a FORTRAN software package called **Multiframe Phase-Diversity Reconstruction** (MPDR), for the purpose of atmospheric image reconstruction using the techniques described in this paper. In the current version (v0.2), the user interface and the atmospheric simulation is based on MATLAB 5.3. For efficiency, the restoration is performed by (compiled) FORTRAN code using MATLAB's Mex interface to interconnect the restoration code and the user interface.

The data dependencies inherent in equations (12) through (16) imply that two passes over the phase data are required. A time/space trade-off exists between storing several fairly large intermediate values for use in the second pass versus re-computing those intermediate values during the second pass. Two versions of the restoration code have been developed. One optimizes for time, but assumes sufficient memory is available. Another optimizes for lower memory use but incurs an increased time cost of six FFTs per function evaluation instead of four. We call it the *small memory* version of MPDR. In both versions, we applied a compiler optimization technique known as *lifetime analysis* to optimize the re-use of temporary arrays.

MPDR uses the highly efficient FFTW (2.1.2) library (see www.fftw.org) for evaluation of the reduced cost functional (12), the gradient (13) and for other computations. Wherever possible, MPDR exploits the well-known Hermitian symmetry property for Fourier transforms of real data. Given $A \in \mathcal{R}^{m \times n}$, let $B = \mathcal{F}(A)$, where \mathcal{F} denotes the 2-D discrete Fourier transform. Then, for $0 \leq i < m$ and $0 \leq j < n$

$$B_{i,j} = \overline{B}_{m-i \text{ modulo } m, n-j \text{ modulo } n} . \quad (20)$$

Note that this symmetry property described in (20) is preserved under addition, conjugation, and component-wise multiplication. MPDR also uses a modified version of the numerical optimization Fortran routine L-BFGS-B from the Argonne National Laboratory Optimization Technology Center to minimize the reduced phase cost functional given in (12). (See the file `deconv.F` in the sub-directory `Restore`.)

Original Matlab m-files in the directory `Simulate` were primarily written by Brent Ellerbroek (AFRL and Gemini) and Curt Vogel (Montana State University), for the purpose of generating test phase screen data allowing for various atmospheric turbulence conditions. Some of the original Matlab m-files in the directory `Utilities` for manipulating the data, etc., were also written by Curt Vogel. In addition, Dave Tyler (University of New Mexico) provided phase screen data generated using FORTRAN code. He and Brent Ellerbroek also gave helpful advice on choosing certain atmospheric turbulence parameters. The directory `Restore` contains FORTRAN routines for the multiframe/multichannel phase-diverse speckle computations. Some of the FORTRAN routines for cost functional and gradient evaluations were translated from Matlab code written by Curt Vogel (MSU) and by Mike Leonard (UCLA). All the production code in MPDR can now be used independent of Matlab.

The basic phase-diverse speckle algorithm (for the special single frame case) which forms the basis of our work is discussed in the paper: C. Vogel, T. Chan and R. Plemmons, "Fast algorithms for phase diversity-based blind deconvolution", in *Adaptive Optical System Technologies*, SPIE Kona Proc. Vol. 3353, pp. 994-105, 1998. (A postscript file containing the paper can be downloaded from the web page: <http://www.wfu.edu/~plemmons/> or from the web pages of T. Chan or C. Vogel). Our current paper represents an extension of that work to multiframe/multichannel phase recovery and object reconstruction.

A compressed tar file of our software can be found at <http://www.mthsc.wfu.edu/~torgerse/mpdr> .

5 Parallel Implementation

As expected with quasi-Newton type methods, it appears there is little opportunity for coarse-grained parallelization in the L-BFGS-B algorithm. Profiling the sequential implementation indicates that the code spends between 90% and 92% of its time evaluating the functional and corresponding gradient in equations

(12) and (13). Thus, some speedup could be hoped for by parallelizing the computation across the frames and channels. However, the parallel computation and distribution of F in equation (16) requires synchronization of all processors and $\mathcal{O}(n^2 \log KT)$ sequential communication at each iteration.¹ This approach is not expected to scale well to large images and a large number of frames.

Despite the skeptical prospects for parallelization described above, a very promising approach to parallelizing the MPDR code can be noted. The key observation is that in simulations of the one frame case, we observe quite good recovery of the phase, but less than desirable restorations of the object image. Thus, it is possible to treat all frames as if they were unrelated, single-frame, two-channel problems in order to recover an approximation $\phi_1, \phi_2, \dots, \phi_T$ to the phases for each frame. The resulting set of T single-frame problems is “embarrassingly parallel.” The approximate single-frame recovered phases $\phi_1, \phi_2, \dots, \phi_T$ do not minimize the functional in equation (12), but they may be used as high-quality starting points for the normal T -frame problem. In practice, it is hoped that a very small number of iterations of the full T -frame problem produces a good restoration. A sequential prototype implementation of the above scheme has been tested and found to yield good restorations using 65 iterations for each individual single-frame problem, and only 3 iterations of the full-sized T -frame problem.

6 Simulation Tests

The numerical experiments we discuss in this section were performed on data produced by simulated multiple time frame, multiple channel phase-diversity data. The phase screens are generated according to a Von Karman turbulence model [9] using a phase screen generator primarily written by Brent Ellerbroek to simulate “seeing conditions” at the Starfire Optical Range. It can be used to generate phase screens for distributed atmospheres, including some or all of the following effects:

- Finite outer scale.
- Time series of phase screens with distinct wind velocities for each layer.
- Anisoplanatism (i.e., phase screens for point sources in different directions).
- Scintillation and diffraction effects.

The principal limitation of the generator is that the phase screens are periodic, so that (a) one cannot expect really long sequences to have the correct temporal statistics, (b) one must keep the aperture diameter less than about 0.5 times the width of the screen for reasonably correct higher-order turbulence statistics, and (c) one should keep the diameter less than about 0.1 times the width of the screen for reasonable tilt-included statistics.

Simulated atmospheric phase was generated according to the multiframe/multichannel phase-diverse speckle model in equation (1). Simulated phase-diversity image data was then generated according to a discretization of the model. Two sample phase screens and their recoveries are shown in Figure 1. Although the pupil may lie within an $n \times n$ grid, a computational grid of size $2n \times 2n$ was used in the Fourier domain to avoid wrap and reduce other edge effects. Discrete Fourier transforms were computed using the 2-D fast Fourier transform package FFTW. To simulate instrument noise, Poisson and Gaussian errors were added to the generated phase data.

To solve the phase recovery and object reconstruction inverse problems in a stable manner, Tikhonov regularization with the regularization functionals described in Section 2.1 were applied. The reduction scheme of Section 2.2 was employed to eliminate the unknown object from the cost functional, thereby reducing the number of unknowns for the optimization routine to Tn^2 point values of the phase, where T is the number of data frames of size $n \times n$. The reduced cost functional was minimized using the implementation of the limited memory BFGS method described in Section 3. Sample true, diffraction limited, blurred and noisy, and reconstructed objects are shown in Figure 2. A table of sample execution times is given in Figure 3. The regularization and stopping parameters were chosen to illustrate typical values and the times shown

¹ $\mathcal{O}(n^2 \log KT)$ refers to problems involving K channels, T frames, and $n \times n$ images.

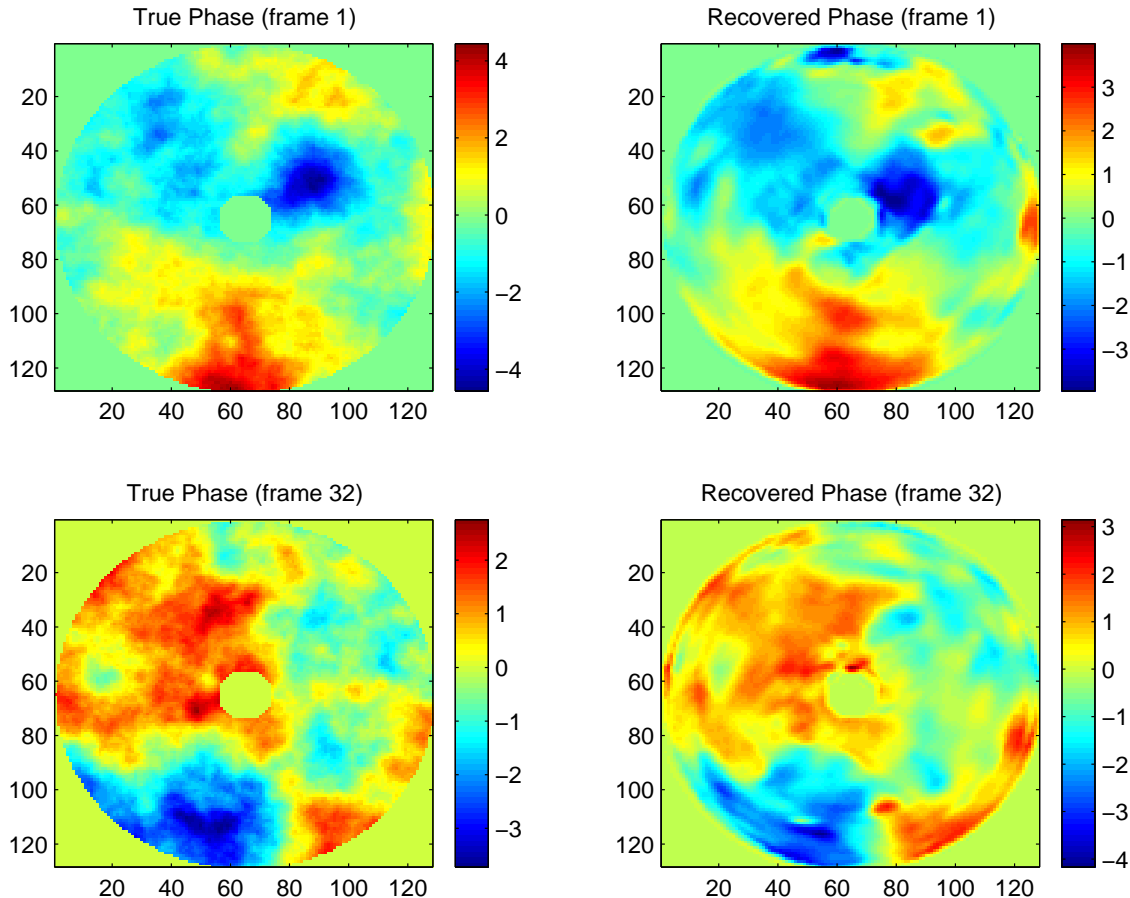


Figure 1: *Sample phase recoveries.* The left-side subplots show the true atmospheric phase screens, or wavefront profiles, and the right-side subplots show the recovered phases using the MPDR software package. For this test 32 frames of 128×128 data were used.

are typical of those needed to achieve an acceptable restoration. This test was performed on a Sun Ultra Enterprise 450 configured with 384MB memory and 250MHz processors.

The performance of the limited memory algorithm depends on the number of vectors s and y saved, as discussed in Section 3. If no vectors are saved, the algorithm reduces to a scaled version of the steepest descent method. On the other hand, if all the vectors are saved, the usual BFGS scheme results. As should be expected, performance improves as more vectors are saved. Significant improvement results in saving 5 vectors as opposed to saving none at all. However, relatively little is gained in going from 5 to more saved vectors, so the algorithm is quite memory and cost efficient.

7 Future Research Directions

Thus far, we have addressed the development of high performance optimization algorithms, and software implementations that are fast and yet storage-efficient on serial computers. Tests on simulated multiple time frame, multiple channel phase-diversity data confirm the effectiveness of our MPDR software package on realistic phase recovery and object reconstruction. However, several tasks remain to be undertaken before this project is fully complete. We hope to consider the following additional work:

- Further investigation of regularization techniques (such as MAP) for the phase, and development of an effective phase unwrapping scheme in the context of wavefront phase recovery.

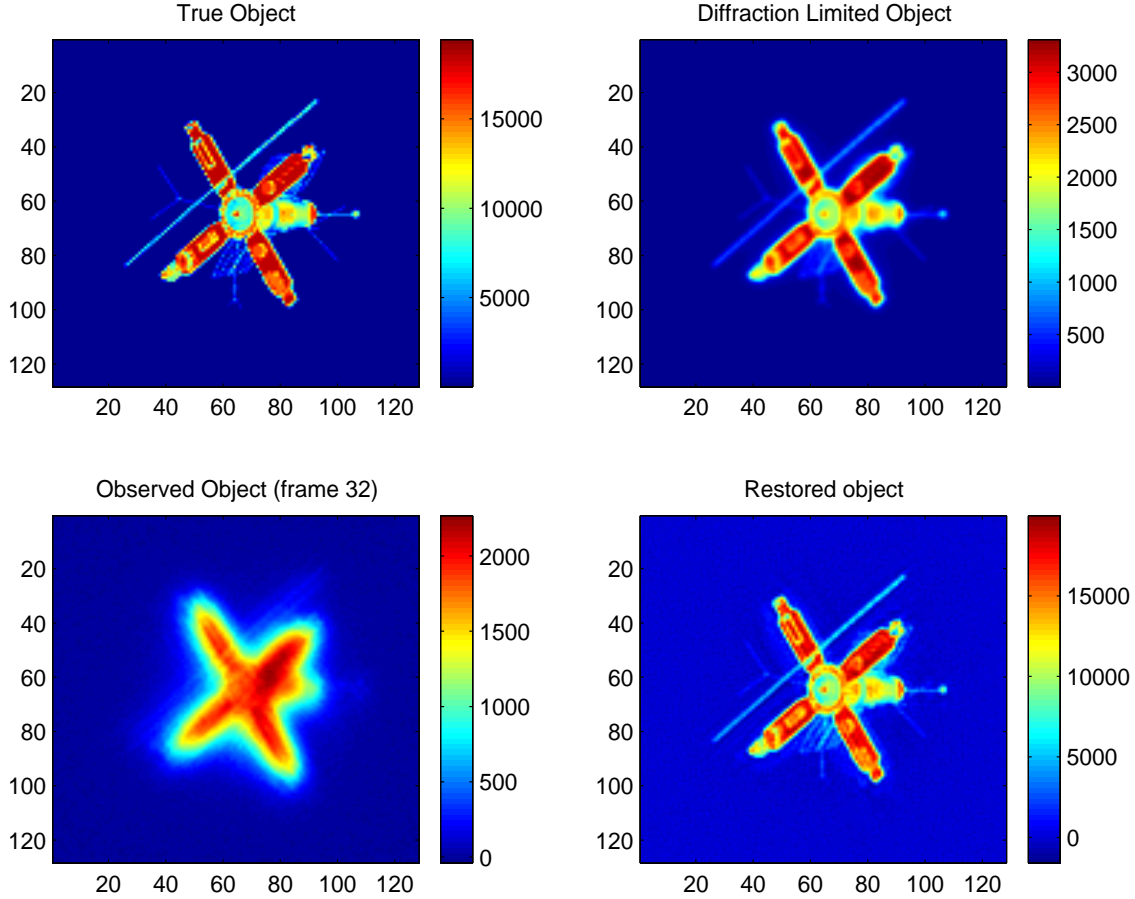


Figure 2: *Sample object recovery using 128×128 resolution. The upper left subplot shows the true object, the upper right subplot shows the true object diffraction limited by the imaging system, the lower left subplot shows the 32nd (last) frame of blurred and noisy data, and the lower right subplot shows the recovered image.*

- Further tests of the algorithms and code using data from AMOS to confirm the effectiveness of our software package on real atmospheric imaging problems.
- Investigation of alternate reduced formulations of the cost functional that avoid the standard Wiener filter restoration step for the image common to PDS methods.
- Parallelization of the codes on the MHPCC IBM SP2 and investigation of the possible incorporation of the MPDR software package into the Maui image manager and on-line systems archive (MIMOSA) at the AMOS facility (see [10]).

We envision that the project described in this paper will result in new technologies for atmospheric image reconstruction, in the form of robust and efficient algorithms as well as their implementation in our MPDR software package. Our techniques will be tested on applications from space object imaging from AMOS. Packaging the results of our research into reliable software will hopefully facilitate the effective and timely transfer of new knowledge to AMOS, as well as to other research laboratories and industry.

Acknowledgments. The authors wish to thank Tony Chan, Brent Ellerbroek, Mike Leonard, Kathy Schulze, Dave Tyler, Curt Vogel and Zac Ziegler for providing very helpful advice during this study. Research by Robert J. Plemmons and Todd C. Torgersen on this project was supported in part by the National Science Foundation under grant CCR-9732070, and research by William K. Cochran was supported by Research Experiences for Undergraduates funds in that grant. This project also received support from the

64 × 64 Problems			128 × 128 Problems		
Frames	Iterations	Elapsed Time	Frames	Iterations	Elapsed Time
1	22	2.9 secs.	1	24	20 secs.
16	36	1.15 min.	16	35	5.5 min.
32	35	2.2 min.	32	36	10.95 min.
64	35	4.4 min.	64	40	30.7 min.

Figure 3: Sample Execution Times with Two Diversity Channels. The small-memory version of the code, used for the 128×128 , 64 frame problem, solves for more than 10^6 phase values at each nonlinear iteration.

HPCERC/Maui Project under AFRL Contract F29601-96-D-0128, Task Order 05, including accounts on the MHPCC IBM SP2 computing system.

References

- [1] R.H. Byrd, P. Lu, J. Nocedal and C. Zhu, “A limited memory algorithm for bound constrained optimization,” *SIAM J. Scientific Computing* **21**, pp. 1190-1208, 1996.
- [2] H. Engl, M. Hanke, and A. Neubauer, *Regularization of Inverse Problems*, Kluwer Academic Publishers, Dordrecht, 1996.
- [3] R.A. Gonsalves, “Phase retrieval and diversity in adaptive optics,” *Optical Eng.* **21**, pp. 829-832, 1982.
- [4] I. Irwan and R. Lane, “Phase retrieval with prior information,” *J. Opt. Soc. Am. A* **15**, pp. 2302-2311, 1998.
- [5] C.T. Kelley, *Iterative Methods for Optimization*, SIAM Press, 1999.
- [6] R.G. Paxman, T.J. Schulz, and J.R. Fineup, “Joint estimation of object and aberrations by using phase-diversity,” *J. Optical Soc. Am., A* **9**, pp. 1072-1085, 1992.
- [7] P. Pauca, B. Ellerbroek, N. Pitsianis, R. Plemmons, and X. Sun. “Performance modeling of adaptive-optics imaging systems using fast Hankel transforms. *Advanced Signal Processing Algorithms, Architectures, and Implementations VIII*, 3461:339–347, 1998.
- [8] P. Pauca and R. Plemmons, “Some computational problems arising in adaptive optics imaging systems,” preprint 1999, to appear in *Computational and Applied Math.*
- [9] M.C. Roggemann and B. Welsh, *Imaging Through Turbulence*, CRC Press, 1996.
- [10] K.J. Schulze, “Maui image manager and on-line systems archive (MIMOSA) – Intelligent observatory data management and processing”, *This Proceedings*, 1999.
- [11] J.H. Seldin, M.F. Reiley, R.G. Paxman, B.E. Stribling, B.L. Ellerbroek, and D.C. Johnston, “Space-object identification using phase-diverse speckle,” in *Digital Image Recovery and Synthesis III, SPIE Proceedings Vol. 3170*, pp. 2-15, 1997.
- [12] D.G. Sheppard, B.R. Hunt, and M.W. Marcellin, “Iterative multiframe superresolution algorithms for atmospheric-turbulence-degraded imagery,” *J. Opt. Soc. Am., A* **15**, pp. 978-992, 1998.
- [13] D.W. Tyler, S.D. Ford, B.R. Hunt, R.G. Paxman, M.C. Roggemann, J.C. Roundtree, T.J. Schulz, K.J. Schulze, D.G. Sheppard, J.H. Seldin, B.E. Stribling, W.C. Van Kampen and B.M. Welsh, “Comparison of image reconstruction algorithms using adaptive optics instrumentation,” in *Adaptive Optical System Technologies, SPIE Kona Proc. Vol. 3353*, pp. 160-171, 1998.
- [14] R.K. Tyson, *Principles of Adaptive Optics, 2nd Edition*, Academic Press, 1998.
- [15] C.R. Vogel, T. Chan, and R.J. Plemmons, “Fast algorithms for phase-diversity-based blind deconvolution”, in *Adaptive Optical System Technologies, SPIE Kona Proc. Vol. 3353*, pp. 994-1005, 1998.