A Multigrid Approach to Two-Dimensional Phase Unwrapping

Gregory Dardyk and Irad Yavneh
Department of Computer Science
Technion

January 12, 2003

Abstract

The two-dimensional phase unwrapping problem is studied. Using the minimum L^p-norm approach, we apply three different nonlinear multigrid algorithms for reconstructing surfaces from their "wrapped" values—two classical approaches and a novel Multilevel Nonlinear Method (MNM). The methods prove to be efficient even for difficult problems with noisy and discontinuous original images. The new method, MNM, exhibits the fastest convergence of the three for all the problems, given an appropriate choice of a damping parameter. Proposed methods for choosing this parameter automatically are mentioned.

1 Introduction

Phase unwrapping is a computational process whereby a surface ϕ (often representing "phase") is reconstructed from its so-called "wrapped" form, ψ. In the absence of noise, $\psi(\mathbf{x})$ is equal to $\phi(\mathbf{x})+2\pi k(\mathbf{x})$, where $k(\mathbf{x})$ is an integer-valued function such that $-\pi < \psi \leq \pi$. The problem of phase unwrapping has drawn considerable interest over the last few years. It is an essential part of many coherent signal processing applications, for example, Interferometric Synthetic Aperture Radar (SAR) and optical interferometers. Coherent processing is based on a single signal property known as phase, which often relates to some physical quantity such as surface topography in interferometry. However, the actual phase values cannot be extracted directly from the physical signal, since phase influences the signal through its principal values that lie between $\pm\pi$ radians. All we can get is the wrapped phase, i.e., the phase values forced into the interval $(-\pi, \pi]$ by a modulo 2π operation.

The unwrapping process aimed at providing an estimation of the actual phase function ϕ, given the wrapped function ψ, turns out to be a difficult problem. To begin with, the relation between ϕ and ψ is nonlinear. Furthermore, interferometric applications deal with real signals in the real world, therefore noise in the measured phase is inevitable, and this complicates the problem appreciably. Phase aliasing is another difficulty that may appear when the sampling resolution is insufficient. Finally, the true solution may exhibit genuine discontinuities which are not easily distinguishable from those produced by the wrapping process.

Many phase unwrapping algorithms were developed during the last ten years. These algorithms may be divided into two general classes: path-following and minimum-norm methods [13]. The path-following, or as they are also called residue-cut "tree" algorithms, unwrap by integrating the gradients (wrapped differences) of the wrapped phase along paths that avoid problematic areas where these differences are inconsistent. The minimum-norm algorithms take a more global approach: they minimize the distance between $\nabla\phi$ and the gradient that is estimated from the wrapped function values. Often, minimum L^p-norm algorithms are used with $p = 2$, yielding a linear Least-Squares problem. Such problems can be solved readily, but the solution is too easily corrupted by noise and data discontinuity (see, e.g., [13] and also examples below). For lower values of p the solution is generally of higher quality but the computational requirements of previously used minimum L^p-norm algorithms are high, due largely to the nonlinearity of the resulting equations. We therefore propose a nonlinear multigrid approach to this problem.

In section 2 we define the discrete minimization problem. In section 3 classical multigrid methods for nonlinear problems are reviewed and the novel MNM algorithm is described. Section 4 demonstrates the results of numerical tests, and conclusions are drawn in section 5.

2 A Minimum-Norm Approach to Two-Dimensional Phase Unwrapping

We describe here the standard minimum-norm approach to estimating the unwrapped phase. We assume that the wrapped phase values are sampled on a rectangular grid, as is typical in practical applications. Accordingly, we denote the wrapped values by $\psi_{i,j}$, and the sought-after unwrapped phase values by $\phi_{i,j}$, with

$$i = 0, 1, 2, \ldots, M, \quad j = 0, 1, 2, \ldots, N. \tag{1}$$

The usual assumption is that if (i, j) and (k, ℓ) are adjacent gridpoints, then $|\phi_{i,j} - \phi_{k,\ell}| < 2\pi$. Accordingly, we define tentative finite-difference approximations to the unwrapped derivatives by

$$\Delta^x_{i,j} = W\{\psi_{i+1,j} - \psi_{i,j}\}, \quad i = 0, \ldots, M-1, \quad j = 0, \ldots, N, \tag{2}$$

$$\Delta^y_{i,j} = W\{\psi_{i,j+1} - \psi_{i,j}\}, \quad i = 0, \ldots, M, \quad j = 0, \ldots, N-1, \tag{3}$$

where W is the wrapping operator defined by

$$W(f_{i,j}) = f_{i,j} + 2\pi k_{i,j}\,, \tag{4}$$

with $k_{i,j}$ chosen such that $W(f_{i,j}) \in (-\pi, \pi]$. Note that when ϕ is in fact discontinuous, $k_{i,j}$ might be different in (2) and (3). Indeed, when this is not the case, the problem is already solved easily by letting $\phi_{i,j} = \psi_{i,j} + 2\pi k_{i,j}$ (and all minimum-norm solutions described below then converge to this solution with zero error).

We require the solution $\phi_{i,j}$ to minimize the discrete functional

$$J^h = \sum_{i=0}^{M-1}\sum_{j=0}^{N} |\phi_{i+1,j} - \phi_{i,j} - \Delta^x_{i,j}|^p + \sum_{i=0}^{M}\sum_{j=0}^{N-1} |\phi_{i,j+1} - \phi_{i,j} - \Delta^y_{i,j}|^p. \tag{5}$$

Differentiating (5) with respect to $\phi_{i,j}$, we obtain the nonlinear system that needs to be solved to obtain the minimum L^p-norm approximation to ϕ:

$$\begin{aligned} (\phi_{i+1,j} - \phi_{i,j} - \Delta^x_{i,j})U(i,j) + (\phi_{i,j+1} - \phi_{i,j} - \Delta^y_{i,j})V(i,j) \\ -(\phi_{i,j} - \phi_{i-1,j} - \Delta^x_{i-1,j})U(i-1,j) - (\phi_{i,j} - \phi_{i,j-1} - \Delta^y_{i,j-1})V(i,j-1) \quad = \quad 0, \end{aligned} \tag{6}$$

where

$$\begin{aligned} U(i,j) &= \begin{cases} |\phi_{i+1,j} - \phi_{i,j} - \Delta^x_{i,j}|^{p-2}, & i = 0, \ldots, M-1,\ j = 0, \ldots, N\,, \\ 0 & \text{otherwise,} \end{cases} \\ V(i,j) &= \begin{cases} |\phi_{i,j+1} - \phi_{i,j} - \Delta^y_{i,j}|^{p-2}, & i = 0, \ldots, M,\ j = 0, \ldots, N-1, \\ 0 & \text{otherwise.} \end{cases} \end{aligned} \tag{7}$$

Note that, for $p < 2$, (6) is a second-order discretization of a nonlinear diffusion equation, whereas for $p = 2$ it reduces to the discrete Poisson equation:

$$(\phi_{i+1,j} - 2\phi_{i,j} + \phi_{i-1,j}) + (\phi_{i,j+1} - 2\phi_{i,j} + \phi_{i,j-1}) = (\Delta^x_{i,j} - \Delta^x_{i-1,j}) + (\Delta^y_{i,j} - \Delta^y_{i,j-1}). \tag{8}$$

This linear equation can be solved easily by multigrid methods as well as others, such as FFT. However, these solutions are often unsatisfactory because of the exaggerated effect of spurious outlying values. Indeed, it is established that better solutions may be found using $p < 2$ (see also section 4). The difficulty is that for such values of p the equations are nonlinear, often with discontinuous and anisotropic coefficients. In the next section we consider three different multigrid algorithms for this problem.

3 Nonlinear Multigrid Methods

Multigrid methods are known for their efficiency in solving large problems arising from the discretization of PDEs, both linear and nonlinear [2, 4, 23]. Linear multigrid algorithms have been considered and applied to the present problem by several authors (see, e.g., [1, 10, 11, 13, 20, 22, 25]). In all these applications the multigrid algorithm is used for solving the Poisson problem that arises from the minimum L^2-norm formulation. As we are interested in lowering p, we consider nonlinear multigrid methods. We next recall two classical approaches, followed by a novel algorithm.

Consider a symmetric nonlinear PDE, discretized on a grid of mesh-size h:

$$N_h(u_h) = f_h\,, \tag{9}$$

where N_h is the discrete nonlinear matrix operator, f_h is a given vector, and u_h is the vector of unknowns. Let v_h denote the current approximation to u_h, and let $e_h = u_h - v_h$ denote the corresponding error. The residual is given by

$$r_h = f_h - N_h(v_h)\,. \tag{10}$$

Substituting $N_h(u_h)$ for f_h in (10) produces

$$N_h(u_h) - N_h(v_h) = N_h(v_h + e_h) - N_h(v_h) = r_h. \tag{11}$$

There are two classical multigrid approaches for solving (11). One method (e.g., [14, 15, 24, 28]), which we refer to as Global Linearization (GL), is to replace the left-hand side by a linear approximation, $K_h e_h$, obtaining,

$$K_h(e_h) = K_h(u_h - v_h) \approx r_h. \tag{12}$$

This equation is solved (approximately) by a linear multigrid solver, and the resulting approximate e_h is added to v_h. This process, which constitutes a single iteration, is repeated until some convergence criterion is satisfied. A common choice is Newton's method:

$$K_h = N_h'[v_h]\,, \tag{13}$$

where $N_h'[v_h]$ is the Jacobian of N_h computed for the current approximation. One drawback of this choice is that K_h is nonsymmetric, which may make (12) relatively difficult to solve. For this reason and others, alternative choices for K_h are often made (see also section 4).

To set the notation for later use, we briefly and informally describe the linear multigrid algorithm. We call grid h the "fine grid", and also define on the same domain a "coarse grid" with mesh-size H. The fine-grid equation (12) is first relaxed by some error-smoothing method, such as Gauss-Seidel relaxation. Then, a coarse-grid problem is defined:

$$K_H(u_H - \bar{I}_h^H v_h) = I_h^H r_h, \tag{14}$$

where K_H is a coarse-grid approximation to K_h, and I_h^H and $\bar{I}_h^H$ are some fine-to-coarse transfer operators (restrictions). After this equation is solved approximately (recursively), the new fine-grid approximation is obtained by interpolating and adding the coarse-grid correction:

$$v_h^{new} = v_h + I_H^h(u_H - \bar{I}_h^H v_h). \tag{15}$$

Here, I_H^h is a coarse-to-fine transfer operator (prolongation). Usually, this is followed by additional relaxation on (12).

A second approach (e.g., [3, 5, 7, 9, 12, 16, 17, 18, 19, 21, 27]), which we refer to as Local Linearization (LL), is the Full Approximation Scheme (FAS) due to Brandt [2]. In this method, the error in the fine-grid solution is smoothed using a nonlinear relaxation method, which employs only a local linearization. Let $\hat{N}_H$ denote a nonlinear coarse-grid approximation to N_h, obtained by rediscretizing the PDE on the coarse grid. The coarse-grid problem is given by

$$\hat{N}_H u_H - \hat{N}_H \bar{I}_h^H v_h = I_h^H r_h. \tag{16}$$

The correction to the fine-grid approximation is added as in (15).

For relatively easy problems, the two approaches usually behave similarly (see for example the discussion in [23]). But for hard nonlinear problems with (nearly) discontinuous nonlinear coefficients, the two approaches generally exhibit distinct behaviors. The GL approach enjoys the advantage that for the linear problems solved in the "inner iterations" there are well-known robust multigrid solvers based on matrix-dependent prolongations and restrictions and Galerkin coarsening. On the other hand, the "outer iteration" converges slowly when the Jacobian is (nearly) singular, and it often requires global methods to approach the basin of attraction. FAS, on the other hand, cannot in general take full advantage of Galerkin coarsening, except at an unacceptable cost. But the use of a direct rediscretization on the coarse grids is not robust, as is well known, and poor behavior is expected (even for linear problems) when the coefficients are discontinuous.

In an effort to combine the advantages of both approaches, we have developed a new Multilevel Nonlinear Method (MNM). We present here a brief description of the method; a detailed presentation appears in [26]. Consider once again the residual equation (11). The main idea of MNM is to split the left-hand side of (11) into a relatively large linear part and a small nonlinear part, and use the (non-robust nonlinear) coarse-grid approximation $\hat{N}_H$ only for the nonlinear part. Accordingly, we add and subtract $K_h(v_h)(u_h - v_h)$ to the left hand of (11), obtaining

$$K_h(v_h)(u_h - v_h) + F[N_h(u_h), N_h(v_h), K_h(v_h), u_h, v_h] = r_h, \tag{17}$$

where

$$F[N_h(u_h), N_h(v_h), K_h(v_h), u_h, v_h] = N_h(u_h) - N_h(v_h) - K_h(u_h - v_h). \tag{18}$$

Note that (17) is equivalent to the original problem, but the nonlinear term F, which is the remainder of a truncated Taylor series, is only of size $O(|u_h - v_h|^2)$, while the first term and the right-hand side are $O(|u_h - v_h|)$. We use the FAS approximation only for the small nonlinear part. This results in the coarse-grid equation,

$$\hat{N}_H(u_H) + (K_H - \hat{K}_H)u_H = I_h^H r_h + \hat{N}_H(\bar{I}_h^H v_h) + (K_H - \hat{K}_H)\bar{I}_h^H v_h\,, \tag{19}$$

where $\hat{K}_H$ is a linear approximation to $\hat{N}_H$ (analogous to the linearization of N_h using K_h). Note that the right-hand terms on both sides of this equation constitute a linear correction to the usual FAS coarse-grid equation, which is intended to compensate for the nonrobustness of FAS. In cases where the linearization yields a poor approximation (and correspondingly the GL approach stagnates), it becomes advantageous to damp these terms (see section 4).

In [26] the MNM method is shown, by analysis and simple numerical examples using exact Newton linearization, to yield a solver that has both the good global convergence properties of FAS and the very fast asymptotic convergence rates typical of Newton iteration. Here, we compare the three approaches for the present difficult and practical problem using a fixed-point linearization.

4 Numerical Tests

We present results of numerical tests and compare the performance of the three proposed algorithms. The tests are performed on a synthetic image of a sine–shaped ramp shown in Fig. 1 (top-left), along with its wrapped form (top-right). The unwrapping of this example is nontrivial because of the data discontinuity: in one direction the wrapped derivatives agree with those of the original image, while in the other direction they agree with those of its wrapped form.

In our implementation we modify slightly the functional which is minimized (following [24]) in order to avoid division by terms that may become vanishingly small. We replace (5) by

$$J^h = \sum_{i=0}^{M-2}\sum_{j=0}^{N-1}(|\phi_{i+1,j} - \phi_{i,j} - \Delta^x_{i,j}|^2 + \beta^2)^{\frac{p}{2}} + \sum_{i=0}^{M-1}\sum_{j=0}^{N-2}(|\phi_{i,j+1} - \phi_{i,j} - \Delta^y_{i,j}|^2 + \beta^2)^{\frac{p}{2}}\,, \tag{20}$$

where β is a small positive parameter. For $\beta \to 0$, (20) reduces to (5), while for $\beta \to \infty$ it reduces to the $p = 2$ functional (multiplied mootly by a constant). We fix $\beta = 0.1$ in all the tests, and check a posteriori that its effect is negligible.

As the robust linear solver (required for the GL and MNM methods) we use Dendy's classical Black-Box multigrid algorithm [6]. For FAS and MNM we need to rediscretize (6) on all the grids. Hence, we must restrict the wrapped phase differences, which we do as follows:

$$(\Delta^x_{i,j})_{2h} = (\Delta^x_{2i,2j})_h + (\Delta^x_{2i+1,2j})_h\,, \quad (\Delta^y_{i,j})_{2h} = (\Delta^x_{2i,2j})_h + (\Delta^x_{2i,2j+1})_h\,. \tag{21}$$

Note that it would be wrong to wrap these coarse-grid gradients.

Since our model problem leads to strongly anisotropic (as well as discontinuous) coefficients, we smooth the errors by a nonlinear alternating line relaxation in red-black ordering (zebra), whereby we first relax at the even-indexed lines and then at the odd-indexed lines. For each line, we need to solve a nonlinear 1D problem which is similar to the full 2D problem. Currently, we use the GL approach for this, repeating the relaxation at lines where a single sweep failed to reduce the residual norm by a factor 0.2 (but giving up after four tries if unsuccessful), and automatically employing under-relaxation if necessary (rarely). Given the demonstrated efficiency of FAS and MNM for the 2D problem, we predict that employing these methods also for the 1D problems in the relaxation will be more efficient, requiring no such reiteration, but this has not yet been implemented.

The finest grid in our tests is 64×64 intervals and the coarsest is 4×4. However, experiments carried out at higher resolutions showed no significant degradation. On the coarsest grid we relax 5 times. We employ V(1,1) cycles with a single alternating zebra sweep (XY) before each coarse-grid visit, and one (YX) sweep after obtaining the coarse-grid correction. For the linearization in the GL approach (and MNM), we use a fixed-point (FP) approach (sometimes called "lagged diffusion"), which has been shown to be more robust than Newton's method for the related problem of total-variation denoising [24]. In this approach, (6) is linearized simply by using ϕ of the previous iteration for U and V. The linearized operator is thus symmetric.

We compare FAS, FP iteration and MNM. Providing a good initial approximation turns out to be important for the convergence of all the algorithms when p is small, therefore we start the solution with $p = 2$ and decrease it by 0.1 steps after every V-cycle until p reaches the desired target value. (These initial iterations are included in the results shown). We expect that an appropriately designed Full Multigrid (FMG) approach will be useful here, but this has not yet been implemented.

In Fig. 1 we show the reconstructed surfaces with $p = 2$ (middle-left), $p = 1.5$ (middle-right), $p = 1$ (bottom-left) and $p = 0$ (bottom-right). It is evident that the solutions improve very significantly as p is decreased. The results for $p = 1$ and for $p = 0$ are already nearly indistinguishable, but $p = 0$ does provide a better solution: the error in this case is $\sim 10^{-6}$ while for $p = 1$ it is $\sim 10^{-3}$. Clearly, the effect of the small but nonzero β is negligible. For this problem, it can be shown that the exact $p = 0$ solution would in fact provide a perfect reconstruction. Of course this is not universally true. For example, if the ramp were much narrower, this method would in fact return the wrapped solution, while for $p = 2$ we would still obtain a "compromise" solution.

In Fig. 2 we compare the residual norm histories for the three algorithms with different values of p. We also include in all the graphs the $p = 2$ histories (denoted LS) for comparison. For MNM, we introduce a damping parameter α for the linear correction terms in (19), obtaining

$$\hat{N}_H(u_H) + \alpha(K_H - \hat{K}_H)u_H = I^H_h r_h + \hat{N}_H(\bar{I}^H_h v_h) + \alpha(K_H - \hat{K}_H)\bar{I}^H_h v_h\,, \tag{22}$$

with $\alpha \in [0, 1]$. Note that $\alpha = 0$ reduces MNM to FAS. For largish p, say 1.8, MNM with $\alpha = 1$ is best. As p is reduced, we find that we need to reduce α in order to maintain fast convergence. In the results shown we use $\alpha = 0.6$ for $p = 1.5$, $\alpha = 0.15$ for $p = 1$, and $\alpha = 0.8$ for $p = 0$. In fact, MNM with $\alpha = 1$ fails to converge for the two small-p cases. This is due to the fact that positivity can be lost for some of the 1D problems encountered in the line relaxation, when the linear correction term is large and happens to be negative. We expect to be able to overcome this problem by an improved relaxation. Also, we plan to develop an automatic selection of α, spatially locally, based on the (local) positivity properties of the operator.

In these tests, FAS converges nearly as fast as MNM, while both these methods seem significantly faster than the GL approach (although at the cost of a computational overhead). However, when we add a Gaussian noise to the image (Fig. 3), MNM proves to be significantly more robust. For $p = 1.5$ and $\alpha = 0.6$, MNM is

Figure 1: Original, wrapped, and reconstructed images with p=2, p=1.5, p=1, and p=0.

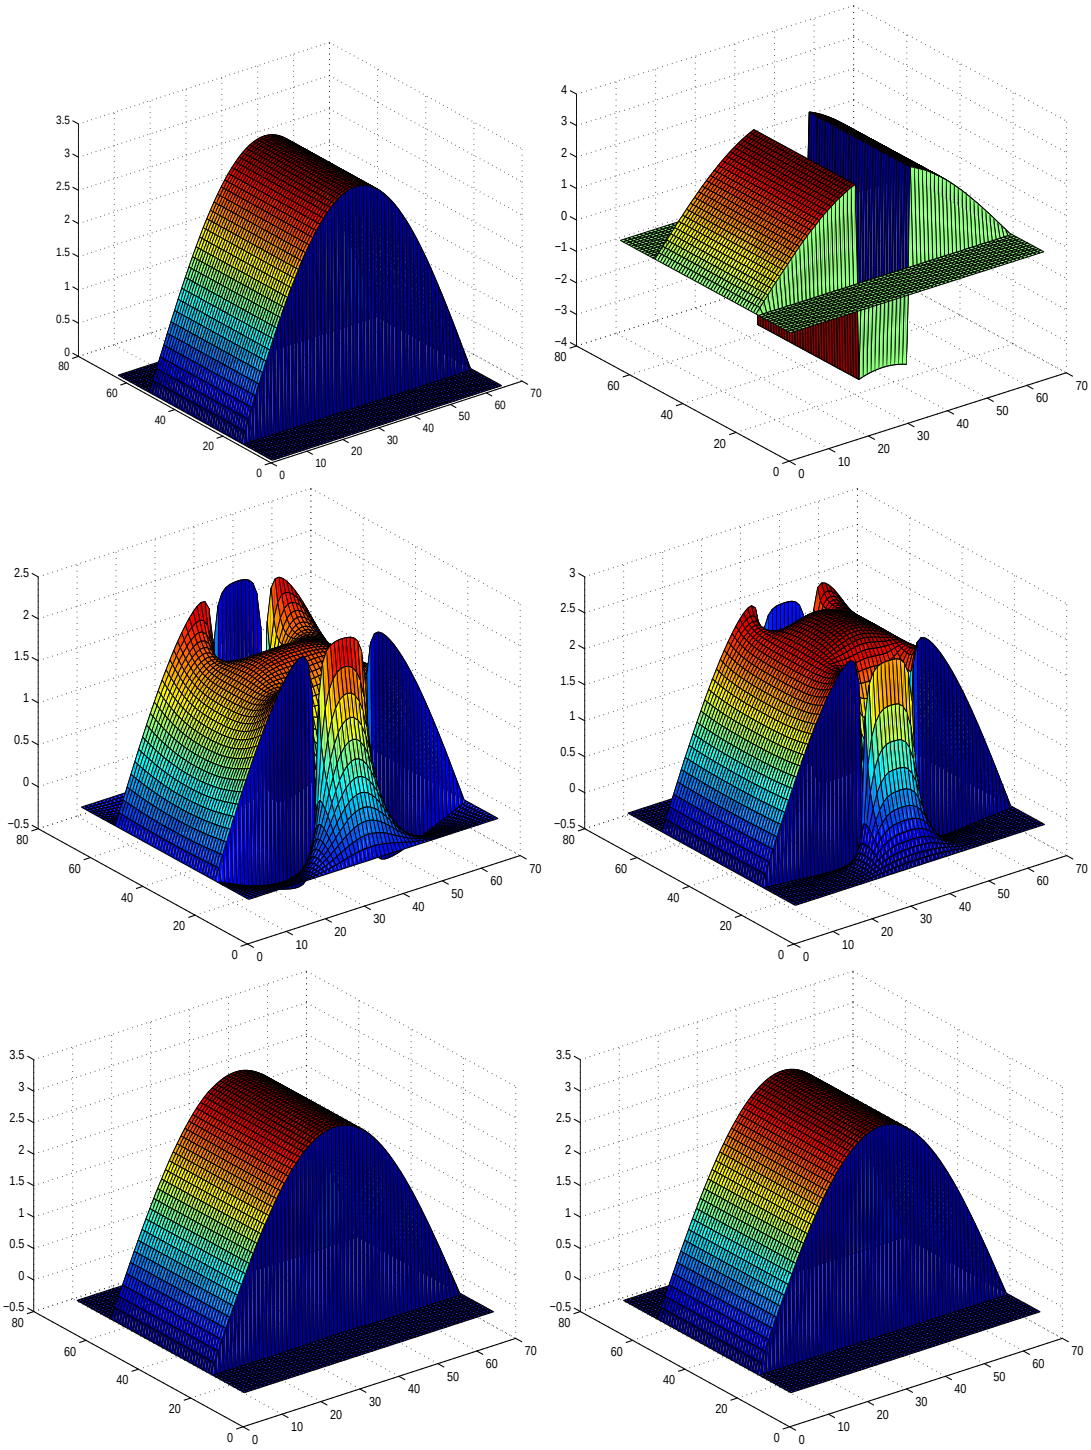

Figure 2: Residual norm convergence histories for p=1.8, p=1.5, p=1, and p=0.

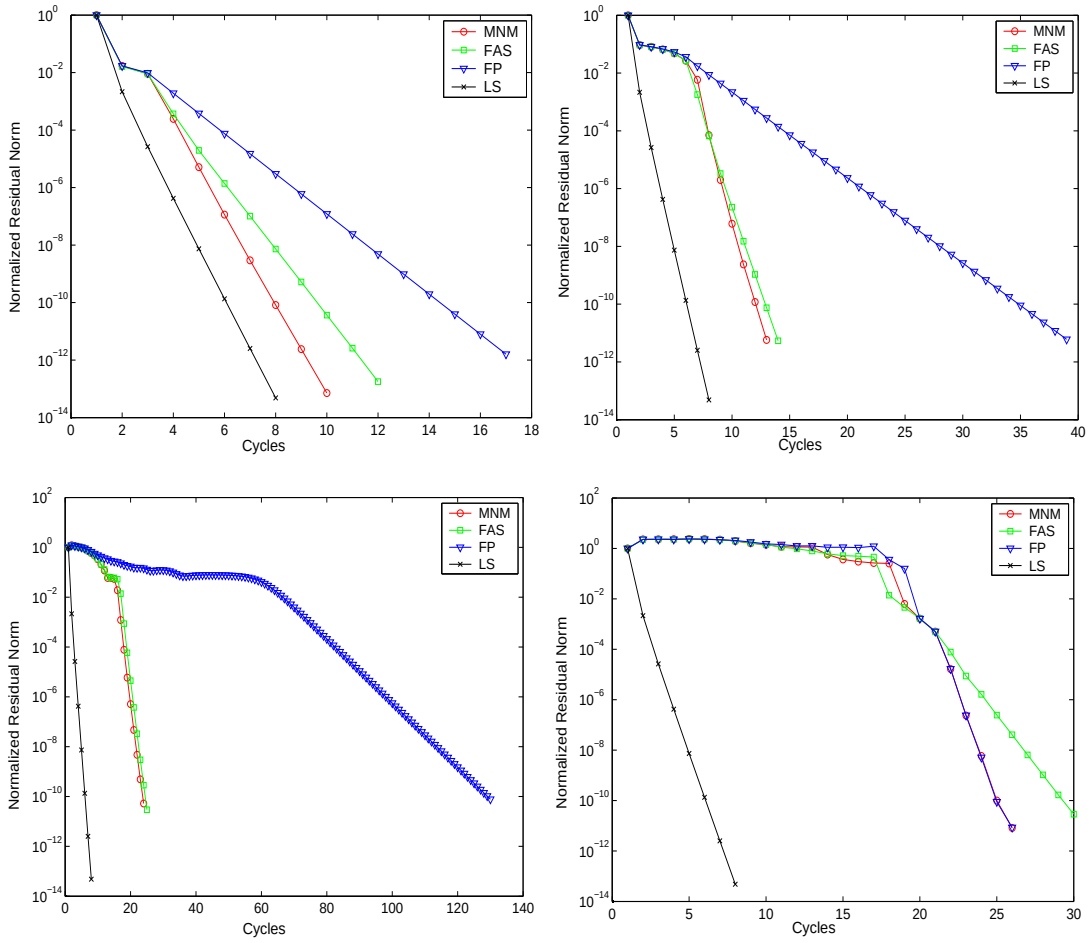

Figure 3: Reconstruction and denoising of a noisy image.

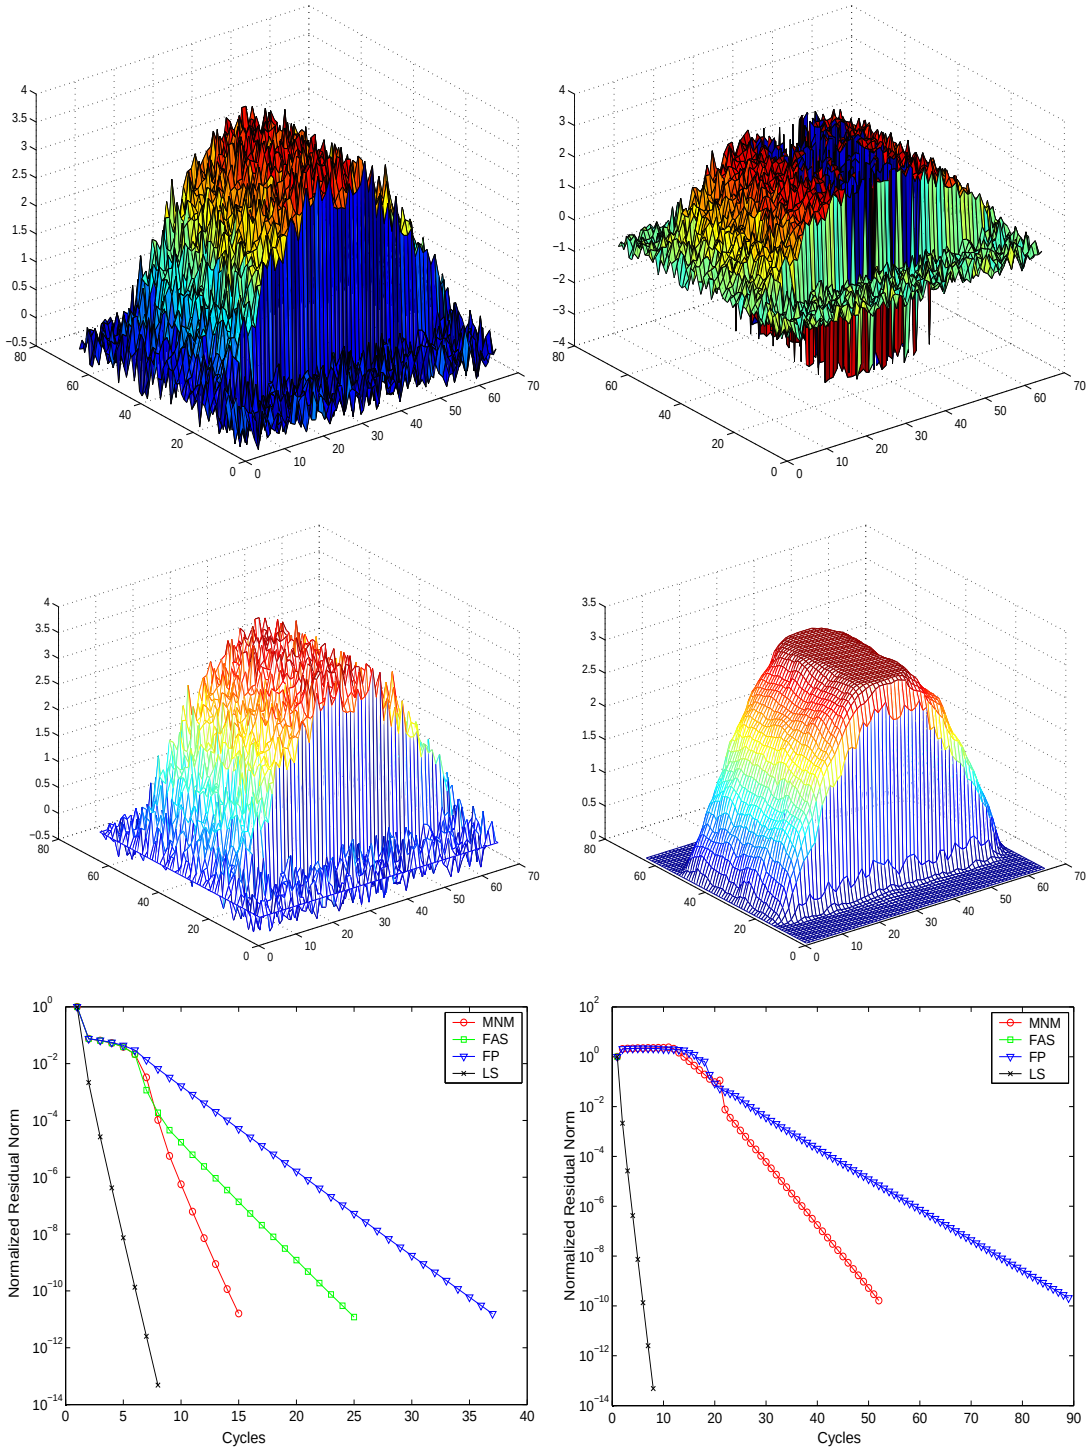

much faster than FAS and FP; and for $p = 0$ and $\alpha = 0.7$, MNM is significantly faster than FP, while FAS fails to converge. The images shown in Fig. 3 are the original noisy image (top-left), it's wrapped form (top right), the $p = 0$ reconstruction (middle-left), and a denoising of the reconstruction using the total variation denoising variant of [24] (middle-right).

5 Conclusions

Nonlinear multigrid methods are shown to be very promising for two-dimensional phase unwrapping. We find that the number of V-cycles required for solving even the L^0-norm minimization problem is not much higher than for the simple Poisson equation of the L^2-norm problem. Compared with this, the execution time of the classical L^p-norm algorithm described in [13] is two orders of magnitude higher than that of linear multigrid. We also find the proposed methods, especially MNM, to handle noisy images well.

Clearly, more tests are required, including real interferometry images. In particular, if strong anisotropies exist which are not aligned with the grid, the basic algorithm may exhibit slower convergence. One effective approach for this problem (which may also improve the present performance), is to apply Krylov subspace acceleration within the nonlinear multigrid cycle using the approach of Oosterlee and Washio [19]. Also, many other algorithmic developments we are currently investigating are expected to improve the overall efficiency appreciably. One natural approach for reducing the overall work is to employ FMG and perform the initial continuation in p on coarse grids. This is expected to eliminate a large part of the discrepancy between the number of fine-grid cycles required for the nonlinear problem, compared to the linear one. Indeed, the excellent asymptotic convergence rates suggest that with a good initial condition only one or two cycles will be required for accurate reconstruction. Improved nonlinear line relaxation methods, based on FAS and MNM are also expected to improve the performance. Finally, an adaptive (spatially local and grid-dependent) choice of the damping α is required to allow MNM to triumph consistently.

Future research will include a combination of unwrapping and total-variation denoising, employing similar nonlinear multigrid methods, to obtain high-quality reconstructions of the original images. We are also interested in extending MNM to unstructured grids. This requires an FAS approach for such problems. Such a method has recently been developed by Dumett et. al. [8], so this is now possible.

References

[1] S. Botello, J. L. Marroquin, and M. Rivera, *Multigrid algorithms for processing fringe-pattern images*, Applied Optics, 37 (1998), pp. 7587–7595.

[2] A. Brandt, *Multi-level adaptive solutions to boundary-value problems*, Mathematics of Computation, 31 (1977), pp. 333–390.

[3] A. Brandt and I. Yavneh, *On multigrid solution of high-reynolds incompressible entering flows*, J. Comput. Phys., 101 (1992), pp. 151–164.

[4] W. L. Briggs, V. E. Henson, and S. F. McCormick, *A multigrid tutorial*, SIAM, Philadelphia, second ed., 2000.

[5] P. M. de Zeeuw, *Nonlinear multigrid applied to a one-dimensional stationary semiconductor model*, SIAM J. Sci. Stat. Comput., 13 (1992), pp. 512–530.

[6] J. E. Dendy (Jr.), *Black box multigrid*, J. Comput. Phys., 48 (1982), pp. 366–386.

[7] D. Drikakis, O. P. Iliev, and D. P. Vassileva, *A nonlinear multigrid method for the three-dimensional incompressible Navier-Stokes equations*, J. Comput. Phys., 146 (1998), pp. 301–321.

[8] M. A. Dumett, P. S. Vassilevski, and C. S. Woodward, *A multigrid method for nonlinear unstructured finite element elliptic equations*, SIAM J. Sci. Comput., (submitted).

[9] B. Epstein, A. Averbuch, and I. Yavneh, *An accurate ENO driven multigrid method applied to 3D turbulent transonic flows*, J. Comput. Phys., 168 (2001), pp. 316–338.

[10] G. Fornaro, G. Franceschetti, R. Lanari, D. Rossi, and M. Tesauro, *Interferometric SAR phase unwrapping using the finite element method*, IEE Proc. – Radar, Sonar Navig., 144 (1997), pp. 266–274.

[11] D. L. Fried, *Adaptive optics wave function reconstruction and phase unwrapping when branch points are present*, Optics Communications, 200 (2001), pp. 43–72.

[12] P. Gerlinger, H. Mobus, and D. Bruggeman, *An implicit multigrid method for turbulent combustion*, J. Comput. Phys., 167 (2001), pp. 247–276.

[13] D. C. Ghiglia and M. D. Pritt, *Two-Dimensional Phase Unwrapping: Theory, Algorithms, and Software*, John Wiley & Sons, Inc., New York, first ed., 1998.

[14] M. J. Holst, R. E. Kozack, F. Saied, and S. Subramaniam, *Treatment of electrostatic effects in proteins – multigrid-based Newton iterative method for solution of the full nonlinear Poisson-Boltzmann equation*, Proteins-Structure Function and Genetics, 18 (1994), pp. 231–245.

[15] M. J. Holst and F. Saied, *Numerical solution of the nonlinear Poisson-Boltzmann equation – developing more robust and efficient methods*, J. Comput. Chem., 16 (1995), pp. 337–364.

[16] M. H. Lallemand and B. Koren, *Iterative defect correction and multigrid accelerated explicit time-stepping schemes for the steady Euler equations*, SIAM J. Sci. Comput., 14 (1993), pp. 953–970.

[17] G. Lonsdale, J. S. Bramley, and D. M. Sloan, *A nonlinear multigrid algorithm and boundary-fitted coordinates for the solution of a two-dimensional flow in a branching channel*, J. Comput. Phys., 78 (1988), pp. 1–14.

[18] J. Molenaar, *Adaptive multigrid applied to a bipolar-transistor problem*, Appl. Numer. Math., 17 (1995), pp. 61–83.

[19] C. W. Oosterlee and T. Washio, *Krylov subspace acceleration of nonlinear multigrid with application to recirculating flow*, SIAM J. Sci. Comput., 21 (2000), pp. 1670–1690.

[20] M. D. Pritt, *Phase unwrapping by means of multigrid techniques for interferometric SAR*, IEEE Trans. Geoscience and Remote Sensing, 34 (1996), pp. 728–738.

[21] J. W. Ruge, Y. Li, S. McCormick, A. Brandt, and J. R. Bates, *A nonlinear multigrid solver for a semi-Lagrangian potential vorticity-based shallow-water model on the sphere*, SIAM J. Sci. Comput., 21 (2000), pp. 2381–2395.

[22] J. Strand and T. Taxt, *Performance evaluation of two-dimensional phase unwrapping algorithms*, Applied Optics, 38 (1999), pp. 4333–4344.

[23] U. Trottenberg, C. Oosterlee, and A. Schüller, *Multigrid*, Academic Press, London and San Diego, 2001.

[24] C. R. Vogel and M. E. Oman, *Iterative methods for total variation denoising*, SIAM J. Sci. Comput., 17 (1996), pp. 227–238.

[25] W. Yao, D. Wu, and A. He, *Real three-dimensional reconstruction of temperature field based on on wavefront retrieval*, Optical Engineering, 37 (1998), pp. 2710–2716.

[26] I. Yavneh and G. Dardyk, *A multilevel nonlinear method*, tech. report, Technion, Haifa 32000, Israel, January 2003.

[27] I. Yavneh and J. C. McWilliams, *Robust multigrid solution of the shallow-water balance equations*, J. Comput. Phys., 119 (1995), pp. 1–25.

[28] S. Zeng and P. Wesseling, *Multigrid solution of the incompressible Navier-Stokes equations in general coordinates*, SIAM J. Numer. Anal., 31 (1994), pp. 1764–1784.