

PROCEEDINGS OF SPIE

SPIDigitalLibrary.org/conference-proceedings-of-spie

Fast algorithms for model-based imaging through turbulence

Sridhar, Venkatesh, Kisner, Sherman, Midkiff, Samuel, Bouman, Charles

Venkatesh Sridhar, Sherman J. Kisner, Samuel P. Midkiff, Charles A. Bouman, "Fast algorithms for model-based imaging through turbulence," Proc. SPIE 11543, Artificial Intelligence and Machine Learning in Defense Applications II, 1154304 (20 September 2020); doi: 10.1117/12.2570789

SPIE.

Event: SPIE Security + Defence, 2020, Online Only

Fast Algorithms for Model-based Imaging through Turbulence

Venkatesh Sridhar ^{*a}, Sherman J. Kisner^b, Samuel P. Midkiff^a, and Charles A. Bouman^a

^aSchool of Electrical and Computer Engineering, Purdue University, West Lafayette, IN

^bHigh Performance Imaging LLC, West Lafayette, IN

ABSTRACT

Digital holography (DH) systems have the potential to perform single-shot imaging through deep turbulence by incorporating emerging algorithms, such as model-based iterative reconstruction (MBIR), that jointly estimate both the phase-errors and speckle-free image. However, the high computational cost of MBIR poses a challenge for use in practical applications.

In this paper, we propose a method that makes MBIR feasible for real-time DH systems. Our method uses surrogate optimization techniques to simplify and speed up the reflectance and phase-error updates in MBIR. Further, our method accelerates computation of the surrogate-updates by leveraging cache-prefetching and SIMD vector processing units on each CPU core. We analyze the convergence and real CPU time of our method using simulated data sets, and demonstrate its dramatic speedup over the original MBIR approach.

Keywords: Digital Holography, MBIR, SIMD parallelism, surrogate optimization, phase-recovery

1. INTRODUCTION

Digital holography (DH) systems can acquire high-resolution images of far-away targets using a coherent laser source and an image-sensor such as a focal plane array (FPA). The main advantage of DH systems is their ability to coherently detect weak signal fields that are modulations of the source optical field. The resulting complex images can be processed to remove severe distortions that would not otherwise be possible with conventional non-coherent detection. Consequently, DH imaging has huge potential for real-time remote-sensing and surveillance (ISR) applications.

Figure 1 illustrates remote-sensing using a DH system. Reflected light from the target is focused onto the FPA using a lens-array. This weak received field is then demodulated by mixing it with a strong reference field that is identical to the source but has a linear pixel-wise phase offset[†]. This demodulation technique is known as *optical heterodyning*^{1,2}, and the resulting FPA measurement is the *hologram*. While most DH systems use simple Fourier-based methods to form the target image from the hologram, advanced image-formation methods^{2,3} can significantly improve quality.

The presence of deep atmospheric turbulence between the target and image sensor can pose a strong challenge to DH systems. As shown in Figure 1, turbulence distorts the point-wise phase of the pupil-plane optical field[‡], or equivalently, the Fresnel diffraction pattern of the target. In this case, DH systems must remove these phase-errors in order to recover a focused image of the target.

Image sharpening (IS)⁴⁻⁶ and model-based iterative reconstruction (MBIR)^{1-3,7} are perhaps the two major classes of algorithms for estimating the unknown phase errors from DH images. The IS and DH-MBIR methods primarily differ in two ways. First, the IS method estimates the complex-valued *reflection coefficient*, which typically has abrupt spatial variations due to *speckle*. Alternatively, DH-MBIR estimates the real-valued *reflectance*^{1,7}, which is typically much smoother as seen in Figure 4. Second, IS and DH-MBIR use very different

^{*}author is currently with Lawrence Livermore National Labs, Livermore, CA. This work was at Purdue University, West Lafayette. Further author information: Send correspondence to Venkatesh Sridhar, E-mail: vsridha@purdue.edu.

[†]projecting the field at an oblique angle onto the FPA produces such a phase-offset

[‡]We assume isoplanatic atmospheric conditions, which allows us to accumulate the phase-errors during wave propagation in the pupil-plane

mathematical frameworks, and consequently, different iterative computations to estimate the phase-errors and the target image.

IS methods are based on maximizing the sharpness of the intensity image specified by the magnitude-squared of reflection coefficient, g . More specifically, IS uses simple Fourier-inversion methods to express g as a function of the detected pupil-plane field, y , and the unknown phase-errors, ϕ . Then, the phase-errors are estimated by maximizing a sharpness metric associated with the image $|g|^2$. However, g is affected by speckle noise[§], which limits the quality of both the sharpened image and the estimated phase-errors¹.

In contrast, DH-MBIR methods use a Bayesian framework to jointly estimate the target reflectance, r , and the unknown phase-errors, ϕ , from the detected pupil-plane field, y . The key advantage of DH-MBIR is that it estimates the reflectance, r , which is much smoother than the reflection coefficient, $|g|^2$, estimated by IS. Therefore, since the unknown has fewer degrees of freedom, the estimation problem can be more accurately solved with with less data.

In this paper, we propose a method to significantly speed up MBIR for real-time DH applications. Our approach uses surrogate optimization to simplify the pixel-wise phase-error and reflectance updates that dominate the computation. We also show how fast parallel SIMD vector processing instructions together with cache prefetching can be used to speed these operations even on a single core of a modern CPU. In our experiments with simulated data sets we verify the convergence of our method and show that we achieve dramatic speedup over the original DH-MBIR approach¹. More specifically, we show that on a single CPU core our method accelerates the reflectance and phase-error updates by a factor of 15.1x and 37.6x respectively as compared to the original approach, and consequently accelerates each DH-MBIR iteration by a factor of 23.7x.

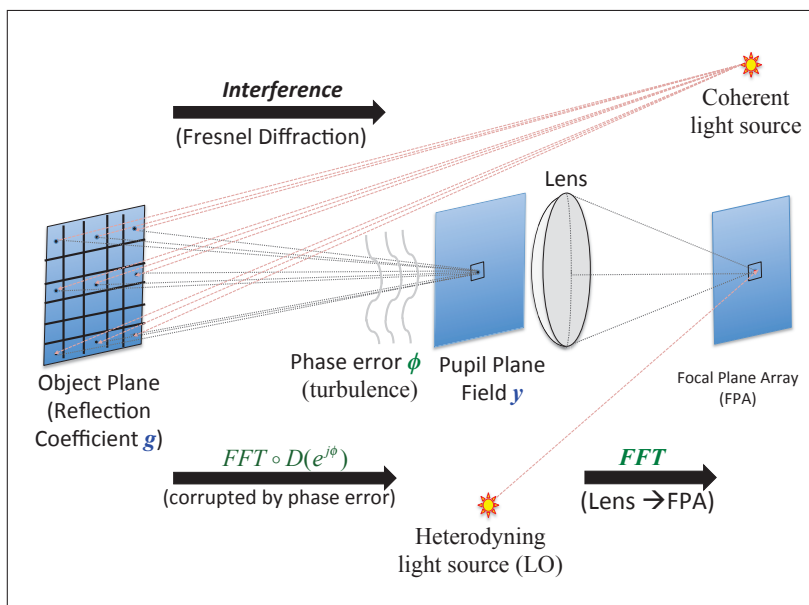


Figure 1. Illustration of a Digital Holography system for remote-sensing. A laser source illuminates the target and a lens focuses the reflected light onto the FPA. The detected field is optically heterodyned with a strong reference field to form a hologram. However, atmospheric turbulence between the target and lens can corrupt the pupil-plane field by inducing phase-errors. These phase-errors must be removed prior to recovering the target image from the hologram.

2. STASTICAL FRAMEWORK FOR DH-MBIR

In this section, we briefly describe the forward model of the DH system and summarize the MBIR reconstruction approach developed by Pellizari et al.¹.

[§] g is speckled since each pixel value can be modeled as a sum of many small scatterers, whose phase is a uniform random variable $\in (0, 2\pi)$

2.1 DH Forward Model

Following from Figure 1, the detected pupil-plane field $y \in \mathbb{C}^n$ is given by

$$y = \mathcal{D}(a)\mathcal{D}(e^{j\phi})Fg + w \quad (1)$$

where $\mathcal{D}(\cdot) = \text{diag}(\cdot)$, $g \in \mathbb{C}^n$ represents the object's reflection coefficient, $F \in \mathbb{C}^{n \times n}$ represents a 2-D Discrete Spatial Fourier transform (DSFT) that computes the Fresnel diffraction integral[¶], $a \in \{0, 1\}^n$ denotes the aperture mask, $\phi \in (-\pi, \pi)^n$ represents the phase-errors in the pupil-plane caused by turbulence, and $w \sim \mathcal{N}(0, \sigma_w^2 I)$ represents white noise. So, we can specify (1) using the following conditional distribution

$$p(y|g, \phi) = \frac{1}{(2\pi\sigma_w^2)^{n/2}} \exp \left\{ -\frac{1}{2\sigma_w^2} \|y - A_\phi g\|^2 \right\}, \quad (2)$$

where $A_\phi = \mathcal{D}(a)\mathcal{D}(e^{j\phi})F$ is the system matrix^{||}.

Since we specify the speckle-free reflectance r as $E[|g|^2]$, we can model g as $\mathcal{N}(0, \mathcal{D}(r))$ shown below

$$p(g|r) = \frac{1}{(\pi^n |\mathcal{D}(r)|)^{1/2}} \exp \{ -g^H \mathcal{D}(r)^{-1} g \}. \quad (3)$$

2.2 MAP Estimation for DH Reconstruction

Pellizari et al.¹ formulate the joint estimation of reflectance $r \in \mathbb{R}^{+n}$ and phase-errors ϕ as a Maximum-a-posteriori (MAP) estimation problem given by

$$(\hat{r}, \hat{\phi}) = \underset{r, \phi}{\text{argmax}} \log p(y|r, \phi) + \log p(r) + \log p(\phi), \quad (4)$$

where $p(y|r, \phi)$ denotes the likelihood model, and, $p(r)$ and $p(\phi)$ denote the prior models for r and ϕ respectively. However, computing the likelihood model in (4) is not tractable without the knowledge of missing information, g . To overcome this issue, Pellizari et al.¹ use the Expectation-Maximization (EM) algorithm to reformulate the MAP estimate of (4) as the following iterative update

$$(r^{(k)}, \phi^{(k)}) = \underset{r, \phi}{\text{argmax}} Q(r, \phi; r^{(k-1)}, \phi^{(k-1)}), \quad (5)$$

where the function Q is specified by

$$Q(r, \phi; r', \phi') = E_g [\log p(y, g|r, \phi)|y, r', \phi'] + \log p(r) + \log p(\phi). \quad (6)$$

In the Appendix, we compute the above Q function starting from (2) and (3). A more comprehensive derivation for the same is available in Pellizari et al.¹.

Algorithm 1 shows the pseudo-code for implementing the EM algorithm specified by (5) and (6) (also see (22) in the Appendix). Each of the M-steps represented by lines 11 and 14 is implemented using only one pass of Iterative Coordinate Descent (ICD) optimization⁸, and the prior models for phase and reflectance are specified by Markov Random Field (MRF) priors.

3. ACCELERATING DH-MBIR

In this section, we propose a method that drastically speeds up the M-steps in Algorithm 1 that dominate the computation by using surrogate optimization and SIMD parallelization.

[¶]we neglect the quadratic phase terms in the Fresnel integral

^{||}As in ¹, we approximate $D(a) = I$, so that A_ϕ is orthogonal

Algorithm 1 EM Algorithm for DH-MBIR

```
1: // Initial values:  $(r', \phi') \in (\mathbb{R}^{+n}, (-\pi, \pi)^n)$ 
2: // Assumption:  $\mathcal{D}(a) = I$ , so  $A_{\phi'}^H A_{\phi'} = nI$ 
3:
4: while not converged do
5:   // E-step: Find Posterior mean  $\mu$  and covariance  $C$ 
6:    $C = \left( \frac{1}{\sigma_w^2} A_{\phi'}^H A_{\phi'} + \mathcal{D}(r')^{-1} \right)^{-1} = \mathcal{D} \left( \frac{r' \sigma_w^2}{r' n + \sigma_w^2} \right)$ 
7:    $\mu = C A_{\phi'}^H y$ 
8:    $z = A_{(\phi=0)} \mu$ 
9:
10:  // M-step: Reflectance update
11:   $r \leftarrow \operatorname{argmin}_{r \in \mathbb{R}^{+n}} \left\{ \sum_{s=1}^n \left( \frac{|\mu_s|^2 + C_{s,s}}{r_s} + \log r_s \right) - \log p(r) \right\}$ 
12:
13:  // M-step: Phase-error update
14:   $\phi \leftarrow \operatorname{argmin}_{\phi} \left\{ -\frac{2}{\sigma_w^2} \operatorname{Real} (y^H \mathcal{D}(e^{j\phi}) z) - \log p(\phi) \right\}$ 
15:   $(r', \phi') \leftarrow (r, \phi)$ 
16: end while
```

3.1 Fast Method for Reflectance Update

Each greedy pixel-wise r update is given by

$$r_s \leftarrow \operatorname{argmin}_{u>0} \left\{ \frac{|\mu_s|^2 + C_s}{u} + \log u + \sum_{j \in \partial s} b_{s,j} \rho(u - r_j) \right\}, \quad (7)$$

where C and μ are defined by lines 6 and 7 of Algorithm 1 respectively, $\rho(\cdot)$ is a Gibbs-prior potential function, more specifically a Q-Generalized Gaussian Markov random field prior (Q-GGMRF)⁹ in this case, and ∂s denotes the neighborhood of pixel s .

Pellizari et al.^{2,3,7} compute the above update almost exactly using a method based on derivative-rooting. However, this approach requires finding the root(s) of a cubic polynomial which is computationally expensive.

In order to speed-up the r update, we replace the objective function in (7) with a suitable surrogate that is easy to minimize. However, designing a surrogate typically requires the objective function to be convex. So first, we substitute the $+\log(\cdot)$ function in (7) with its 1st order approximation, which yields a convex optimization of the form

$$\hat{r}_s \leftarrow \operatorname{argmin}_{u>0} \left\{ \frac{|\mu_s|^2 + C_s}{u} + \frac{u}{r_s} + \sum_{j \in \partial s} b_{s,j} \rho(u - r_j) \right\}. \quad (8)$$

We then design a quadratic surrogate function for (8) based on the Linear Interpolation of Derivative (LID) method¹⁰. We define the function $l(\cdot) : \mathbb{R}^+ \rightarrow \mathbb{R}$ as

$$l(u) = \frac{|\mu_s|^2 + C_s}{u} + \frac{u}{r_s}. \quad (9)$$

The first and second derivatives of $l(\cdot)$ are given by

$$l'(u) = -\frac{(|\mu_s|^2 + C_s)}{u^2} + \frac{1}{r_s}, \quad l''(u) = 2\frac{(|\mu_s|^2 + C_s)}{u^3}.$$

In accordance with the LID method, we specify a quadratic surrogate function that matches the derivative of $l(\cdot)$ defined by (9) at 2 different points, r_s and α_s , where r_s is the current value of the pixel being updated and α_s is another suitably chosen value in the search direction. In this case, our reflectance update is given by

$$\hat{r}_s \leftarrow \operatorname{argmin}_{u \in \mathcal{S}} \left\{ \frac{\theta_2}{2} u^2 + \theta_1 u + \sum_{j \in \partial s} b_{s,j} \rho(u - r_j) \right\}, \quad (10)$$

where \mathcal{S} denotes the search interval $(\min(r_s, \alpha_s), \max(r_s, \alpha_s))$, and, θ_2, θ_1 are given by

$$\begin{aligned} \theta_2 &= \frac{l'(r_s) - l'(\alpha_s)}{r_s - \alpha_s} = (|\mu_s|^2 + C_s) \left(\frac{r_s + \alpha_s}{r_s^2 \alpha_s^2} \right) \\ \theta_1 &= l'(r_s) - \theta_2 r_s = -\frac{(|\mu_s|^2 + C_s)}{r_s^2} + \frac{1}{r_s} - \theta_2 r_s. \end{aligned}$$

We select α_s that bounds our search interval as shown below

$$\alpha_s = \begin{cases} \frac{r_s}{1+M} & \text{if } l'(r_s) + \sum_{j \in \partial s} b_{s,j} \rho'(r_s - r_j) > 0 \\ r_s(1+M) & \text{else,} \end{cases}$$

where $M > 0$. In practice, we suitably vary M such that it decays slowly with the number of iterations.

We can replace the symmetric Gibbs-prior model in (10) with a quadratic surrogate function of its own. Consequently, (10) further simplifies to

$$\hat{r}_s \leftarrow \operatorname{argmin}_{u \in \mathcal{S}} \left\{ \frac{\theta_2}{2} u^2 + \theta_1 u + \sum_{j \in \partial s} b_{s,j} \frac{\tau_j}{2} (u - r_j)^2 \right\}, \quad (11)$$

where τ_j is given by

$$\tau_j = \begin{cases} \frac{\rho'(r_s - r_j)}{(r_s - r_j)} & r_s \neq r_j \\ \rho''(0) & r_s = r_j. \end{cases}$$

So, the surrogate-based reflectance update is given by

$$\hat{r}_s \leftarrow \operatorname{clip} \left\{ -\left(\frac{\theta_1 - \sum_{j \in \partial s} b_{s,j} \tau_j r_j}{\theta_2 + \sum_{j \in \partial s} b_{s,j} \tau_j} \right), \mathcal{S} \right\}, \quad (12)$$

where the $\operatorname{clip}\{x, (a, b)\}$ here denotes x when $x \in (a, b)$, a when $x < a$, and b otherwise.

3.2 Fast Method for Phase-error Update

The exact pixel-wise phase-error update is given by

$$\phi_s \leftarrow \operatorname{argmin}_{u \in [-\pi, \pi]} \left\{ -m_s \cos(u - \varphi_s) + \sum_{j \in \partial s} b_{s,j} \rho(u - \phi_j) \right\}, \quad (13)$$

where $\chi_s = 2y_s z_s^* / \sigma_w^2$ (see line 8 of Algorithm 1 for definition of z), $m_s = |\chi_s|$, $\varphi_s = \angle \chi_s$, and $\rho(\cdot)$ is the phase-wrapped quadratic error with the form

$$\begin{aligned} \rho(\Delta) &= ([(\Delta + \pi) \bmod 2\pi] - \pi)^2 \\ &= \min_n (\Delta + 2\pi n)^2. \end{aligned} \quad (14)$$

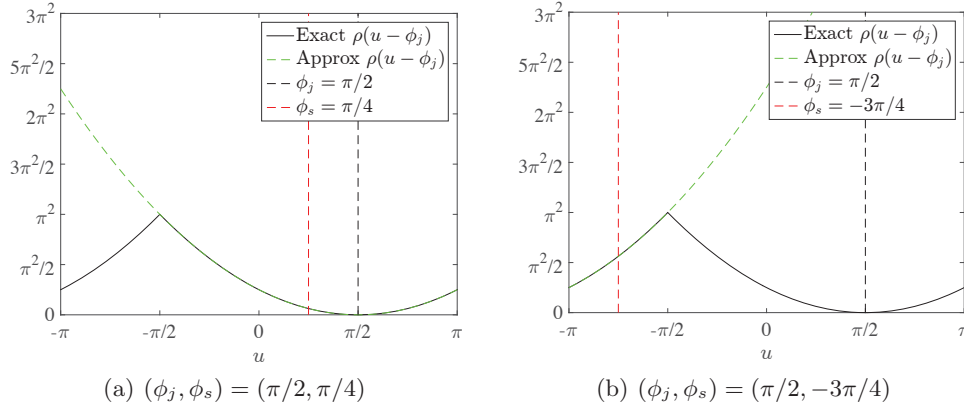


Figure 2. Approximation of the wrapped-phase penalty function $\rho(u - \phi_j)$, $j \in \partial s$, for the s -th pixel ICD update. Subfigures (a) and (b) show how the approximation varies with different values of (ϕ_j, ϕ_s) .

Pellizari et al.^{1,2} estimate the phase-errors on a lower-resolution grid and compute the above ICD update using the Golden-section search (GSS) method. However, the GSS method is computationally expensive since it involves nested evaluations of the local cost function.

We devise a faster update strategy for the ICD update. We first approximate the $\cos(\cdot)$ function in (13) with its Taylor's series expansion of $\cos(\Delta) \approx 1 - \frac{1}{2}\Delta^2$ to yield approximate ICD update of

$$\phi_s \leftarrow \operatorname{argmin}_{u \in [-\pi, \pi]} \left\{ \frac{1}{2} m_s (u - \varphi_s)^2 + \sum_{j \in \partial s} b_{s,j} \rho(u - \phi_j) \right\}. \quad (15)$$

Substituting in the form of $\rho(\cdot)$ from (14), yields the update

$$\phi_s \leftarrow \operatorname{argmin}_{u \in [-\pi, \pi]} \left\{ \frac{1}{2} m_s (u - \varphi_s)^2 + \sum_{j \in \partial s} b_{s,j} \min_{n_j} (u - \phi_j + 2\pi n_j)^2 \right\}. \quad (16)$$

While exact solution of equation (16) is difficult, we can get a local minimum using alternating minimization. One iteration of alternating minimization is given by

$$n_j \leftarrow \operatorname{argmin}_{m \in \mathcal{Z}} (\phi_s - \phi_j + 2\pi m)^2, \quad j \in \partial s \quad (17)$$

$$\phi_s \leftarrow \operatorname{argmin}_{u \in [-\pi, \pi]} \left\{ \frac{m_s}{2} (u - \varphi_s)^2 + \sum_{j \in \partial s} b_{s,j} (u - \phi_j + 2\pi n_j)^2 \right\}. \quad (18)$$

Figure 2 illustrates how the above method approximates the wrapped-phase penalty function in (15). Note that since $\phi_j, \phi_s \in [-\pi, \pi)$, we can show that (17) specifies $n_j \in \{-1, 0, 1\}$, and further, $\phi_j - 2\pi n_j \in (-2\pi, 2\pi)$. Consequently, equation (18) can be simplified further as

$$\phi_s \leftarrow \tilde{\phi}_s + 2\pi k_s \quad (19)$$

where $\tilde{\phi}_s$ is defined as

$$\tilde{\phi}_s = \frac{m_s \varphi_s + 2 \sum_{j \in \partial s} b_{s,j} (\phi_j - 2\pi n_j)}{m_s + 2 \sum_{j \in \partial s} b_{s,j}},$$

and k_s is defined as

$$k_s = \begin{cases} 1 & \tilde{\phi}_s < -\pi \\ 0 & \tilde{\phi}_s \in [-\pi, \pi) \\ -1 & \tilde{\phi}_s > \pi. \end{cases}$$

3.3 SIMD Parallelization of ICD updates

ICD optimization is a serial pixel-wise greedy minimization method. However, from (7) and (13) we notice that the ICD updates for those pixels that do not form pair-wise MRF cliques with one another are independent and can be computed simultaneously. For example, Figure 3(a) shows that in the case of a MRF with symmetric 8-point neighborhood, we can specify a tiled pixel-grid of 4 different colors where ICD updates for pixels of the same color are fully disassociated. Furthermore, the simplified ICD updates represented by (12) and (19) can be fully implemented by basic addition and multiply operations**. So, we can utilize SIMD units that perform fast vector addition and multiplication on a single CPU core to update multiple disassociated pixels at a time. Figure 3(b) illustrates our idea of ICD parallelization using SIMD processing. Importantly, note that the required data must be packed into contiguous arrays prior to SIMD computation.

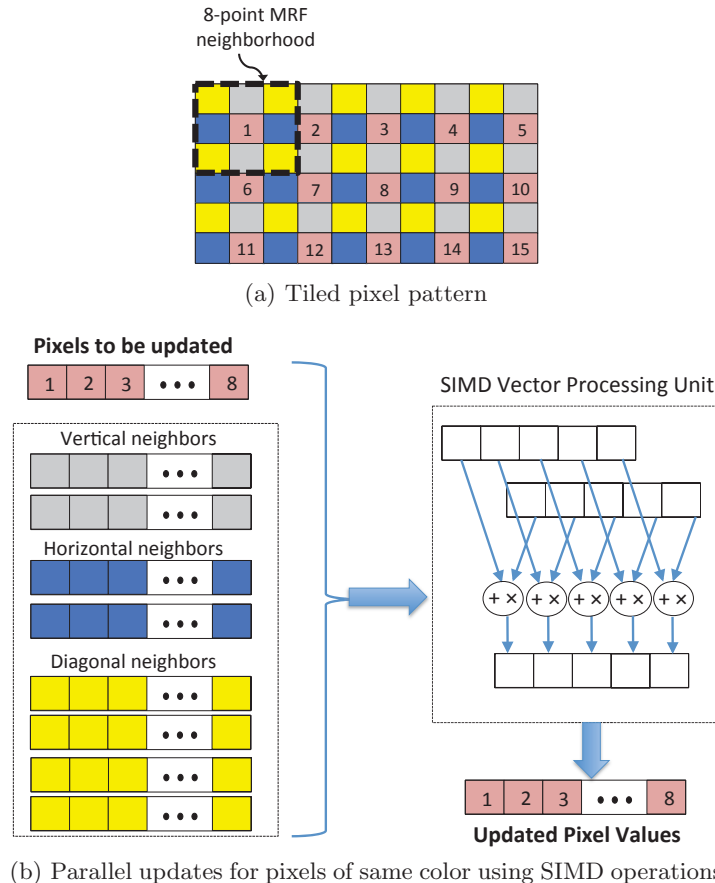


Figure 3. SIMD Parallelization of ICD updates: (a) A tiled grid where pixels of the same color have independent ICD updates (b) SIMD processing units of a CPU can be used to update multiple pixels of the same color at a time using fast vector add and multiply operations.

However, our above approach has one key limitation. Accessing the data in an interleaved manner adds significant overhead that reduces the benefit of the SIMD parallelism. To overcome this issue, we provide an alternate method that negates the need to sample data in a tiled fashion.

We alternatively propose a simple row-wise update strategy which extends the SIMD vector processing method of Figure 3(b) towards simultaneously computing the ICD updates for pixels in contiguous blocks within the same row. In order to explain why this method works in practice, let us assume that within the symmetric 8-point MRF neighborhood depicted in Figure 3(a), the contribution of the horizontal left neighbor is small

**We choose Q-GGMRF parameters⁹ in (7) as $(q, p) = (2, 1)$ or $(2, 2)$ so that its surrogate coefficients τ_j in (11) can be computed easily

compared to the other 7 neighbors put together. In this case, we can show that our row-wise SIMD method approximates serial pixel-wise minimization in raster order.

So, our row-wise update strategy provides dual benefits of fast memory access and SIMD parallelism. Our experiments demonstrate that this method significantly speeds-up ICD computation without affecting convergence.

4. EXPERIMENTAL RESULTS

4.1 Method

In this section, we compare the computational performance, convergence and reconstruction quality of our fast DH-MBIR method against the exact DH-MBIR method¹ using simulated datasets. We implement both MBIR methods on an Intel Xeon CPU (E5-2660 v3) in ANSI C with single-threading. The source code was compiled with the Intel *icc* compiler (ver. 17.0.1) which automatically converts the tight vectorizable loops into AVX2 instructions for SIMD processing. We implement all Fourier-based operations in MBIR using FFT routines from the Intel Math Kernel Library (MKL).

For our experiments, we simulate 6 different data sets, each based on the same object reflectance shown in Figure 4 but a different phase-error screen and reflection-coefficient.

We simulate the phase-errors in the pupil-plane based on a Power Spectral density (PSD) modeling technique^{11,12}. For turbulent conditions, the 2-D spatial distribution of the phase-errors can be modeled as a *Kolmogorov* PSD in the Fourier domain (see equation (1) of Srinath et al.¹²). So, we generate the phase-errors by first scaling white noise in the Fourier domain with a Kolmogorov PSD and then applying an inverse FFT^{1,12}. It is worth noting that the PSD incorporates a key parameter known as the *Fried* coherence length, r_0 , that determines the spatial correlation of the phase-errors. We specifically parameterize half of our simulated phase-error screens by $D_a/r_0 = 10$, and the other half by $D_a/r_0 = 20$, where D_a denotes aperture diameter. The top row of Figure 6(c) and (d) illustrates the simulated phase-screens.

We generate the reflection-coefficient based on the given object reflectance by sampling from the i.i.d complex Gaussian distribution specified by (3). Subsequently, we simulate the detected field in the pupil-plane based on the forward model specified by (1).

For most experiments, the grid-size for both the pupil-plane detection and the reconstruction is 256×256 . For CPU timing experiments alone we use a grid size of 128×128 .

We quantify the quality of our phase-error estimate based on the Strehl ratio metric specified by

$$\text{Strehl ratio} = \frac{\iint_{\text{pupil}} a(x, y) \exp \left\{ j(\phi(x, y) - \hat{\phi}(x, y)) \right\} dx dy}{\iint_{\text{pupil}} a(x, y) dx dy},$$

where a represents the aperture mask and $\phi - \hat{\phi}$ represents the estimation error.

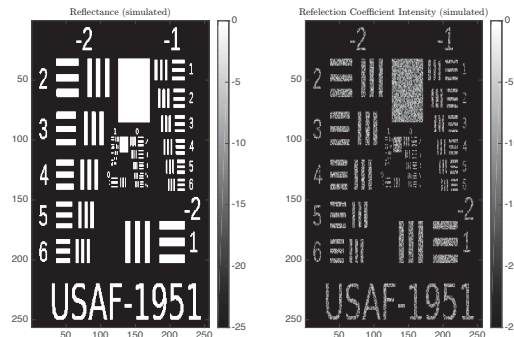


Figure 4. Example depicting object reflectance, r and simulated reflection coefficient, g . Unlike $|g|^2$ (right), the reflectance r (left) is speckle-free and smooth.

Table 1. CPU Time per EM iteration

DH-MBIR Method	E-step	Reflectance update ¹		Phase-error update ²	Full Iteration ³
		$(q, p) = (2, 1)$	$(q, p) = (2, 2)$		
Exact	0.13 ms	1.25 ms	0.80 ms	9.02 ms	9.95 ms
Fast (No SIMD)	0.13 ms	0.53 ms	0.33 ms	0.38 ms	0.84 ms
Fast (with SIMD)	0.13 ms	0.10 ms	0.05 ms	0.24 ms	0.42 ms


¹ Prior model is a Q-GGMRF with 4-point neighborhood and power-parameters specified by (q, p)

² Prior model is a wrapped-phase GMRF with 8-point neighborhood


³ full iteration time calculated for case $(q, p) = (2, 2)$

Table 2. Speedup of Fast DH-MBIR over Exact DH-MBIR

Reflectance update		Phase-error update	Full Iteration
$(q, p) = (2, 1)$	$(q, p) = (2, 2)$		
12.50	15.09	37.58	23.7



(a) 3 different datasets simulated with $D_a/r_0 = 10$

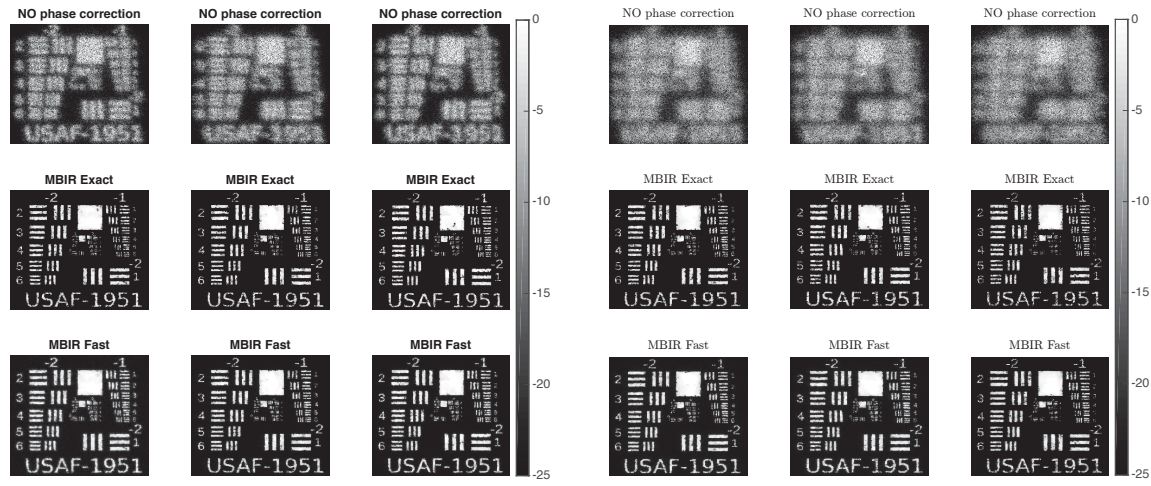


(b) 3 different datasets simulated with $D_a/r_0 = 20$

Figure 5. Effect of fast updates for estimating reflectance, r , and phase-error, ϕ , on the convergence rate of DH-MBIR. The number of iterations for DH-MBIR convergence is unchanged when we use fast updates for ϕ estimation instead of exact updates, but increases when we use fast updates for r estimation in place of exact updates.

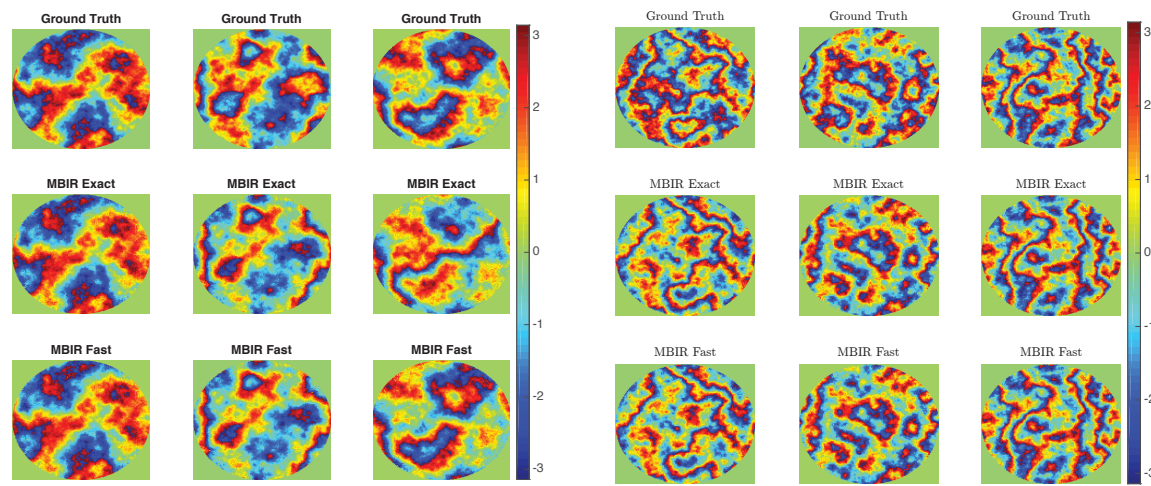
4.2 Results

Table 1 shows the per-iteration CPU time for the exact and fast DH-MBIR methods. Notice that even without utilizing SIMD processing, the surrogate ICD updates are significantly faster than the exact ICD updates. In particular, we observe more than a 20x reduction in time for the phase-error estimation that dominates the exact DH-MBIR method. Furthermore, notice that SIMD parallelization provides additional acceleration of the surrogate ICD updates. Specifically, this speedup resulting from SIMD is more prominent in the case of



(a) Reflectance estimate, $D_a/r_0 = 10$

(b) Reflectance estimate, $D_a/r_0 = 20$



(c) Phase-errors estimate, $D_a/r_0 = 10$

(d) Phase-errors estimate, $D_a/r_0 = 20$

Figure 6. Comparison of reconstruction quality for 6 different simulated datasets. For each dataset, the reflectance estimate without any phase recovery is shown in the top row of sub-figures (a) and (b) respectively, and the corresponding phase-errors in the pupil-plane are shown in the top row of sub-figures (c) and (d). In all sub-figures, the middle row shows reconstruction using the exact MBIR approach while the bottom rows shows reconstruction using the fast MBIR approach. For most datasets, the reconstructions using the fast MBIR method are almost indistinguishable from the exact MBIR method.

reflectance estimation as compared to phase-error estimation^{††}.

Table 2 shows the speedup of our fast DH-MBIR method over the exact DH-MBIR method based on the CPU time in Table 1. The combined effect of surrogate optimization and SIMD parallelization provides a speedup of 15.1x and 37.6x for the reflectance and phase-error updates respectively, and consequently accelerates each DH-MBIR iteration by a factor of 23.7x.

Figure 5 shows the effect of fast updates on the convergence rate of DH-MBIR. Notice that replacing the exact updates for phase-error estimation with the fast updates does not significantly alter the convergence rate. However, when we use fast updates for reflectance estimation in place of the exact updates, we require more iterations to converge.

^{††}the timing includes complex-arithmetic computation of parameters m_s and ϕ_s in (13) which is intensive and not SIMD-compatible. Further using SIMD processing to compute (17) does not provide any significant benefit.

Figure 6 compares the reconstruction quality of the fast and exact DH-MBIR approaches. We can see that for most datasets, the reflectance and phase-error estimates are almost indistinguishable. For the third dataset alone in Figure 6(c), the phase-error estimate from the fast DH-MBIR approach is more accurate as compared to the exact DH-MBIR method.

5. CONCLUSION

In this paper, we proposed a method that drastically reduces the computational cost DH-MBIR for real-time coherent imaging through turbulence. DH-MBIR jointly estimates the speckle-free object reflectance and atmospheric phase-errors from holographic sensor measurements. In order to speedup DH-MBIR, we first designed a surrogate function for the reflectance updates and a simple alternating minimization scheme for the wrapped phase-error updates. Further, we introduced a scheme that accelerates computation of the above surrogate updates using SIMD vector processing instructions. We demonstrated the effectiveness of our fast DH-MBIR method for real-time reconstruction with simulated data sets.

A. APPENDIX

To compute the Q-function in (6), we first need to compute $\log p(y, g|r, \phi)$ as well as the posterior distribution $p(g|y, r', \phi')$. The former is specified by

$$\begin{aligned} \log p(y, g|r, \phi) &= \log p(y|g, \phi) + \log p(g|r) \\ &= -\frac{1}{2\sigma_w^2} \|y - A_\phi g\|^2 - g^H \mathcal{D}(r)^{-1} g - \log |\mathcal{D}(r)| + \text{const.} \end{aligned} \quad (20)$$

Similarly, we can obtain the posterior distribution as

$$\begin{aligned} p(g|y, r', \phi') &= p(y, g|r', \phi') / p(y|r', \phi') \\ &= \frac{1}{z} \exp \left\{ -\frac{1}{2\sigma_w^2} \|y - A_{\phi'} g\|^2 - g^H \mathcal{D}(r')^{-1} g \right\} \end{aligned}$$

where z is a normalizing constant. The above can be more compactly expressed as a complex Gaussian distribution

$$p(g|y, r', \phi') = \frac{1}{z} \exp \{ -(g - \mu)^H C^{-1} (g - \mu) \},$$

where mean μ and covariance matrix C are given by

$$\begin{aligned} C &= \left(\frac{1}{\sigma_w^2} A_{\phi'}^H A_{\phi'} + \mathcal{D}(r')^{-1} \right)^{-1} \\ \mu &= C A_{\phi'}^H g \end{aligned}$$

Consequently the posterior mean and variance are given by

$$E[g|r', \phi'] = \mu \quad \text{and} \quad E[gg^H|r', \phi'] = \mu\mu^H + C. \quad (21)$$

If we assume $\mathcal{D}(a) = I$ in (1), then A_ϕ is orthogonal, and so C is diagonal. Then from (20) and (21), we can show that

$$E[\log(y, g|r, \phi)|r', \phi'] = -\sum_{s=1}^n \left(\frac{|\mu_s|^2 + C_{s,s}}{r_s} + \log r_s \right) + \frac{2}{\sigma_w^2} \text{Real}(y^H \mathcal{D}(e^{j\phi}) A_0 \mu) + \text{const.} \quad (22)$$

The full derivation is available in¹, Appendix B.

6. ACKNOWLEDGEMENTS

This material is based upon work supported by the United States Air Force under contract number FA9451-18-P-0250. We also thank Mark Spencer, Air Force Research Laboratory and Casey Pellizari, US Air Force Academy for their expertise and valuable feedback regarding this work.

REFERENCES

- [1] Pellizari, C. J., Spencer, M. F., and Bouman, C. A., “Phase-error estimation and image reconstruction from digital-holography data using a Bayesian framework,” *Journal of the Optical Society of America A* **34**(9), 1659–1669 (2017).
- [2] Pellizari, C. J., Spencer, M. F., and Bouman, C. A., “Optically coherent image reconstruction in the presence of phase errors using advanced-prior models,” *Proceedings of SPIE, Defense + Security Conference* **10650** (2018).
- [3] Pellizari, C. J., Trahan III, R., Zhou, H., Williams, S., Williams, S. E., Nemati, B., Shao, M., and Bouman, C. A., “Optically coherent image formation and denoising using a plug and play inversion framework,” *Adaptive Optics* **56**(16) (2017).
- [4] Thurman, S. T. and Fienup, J. R., “Phase-error correction in digital holography,” *Journal of the Optical Society of America A* **25**(4), 983–994 (2008).
- [5] Sulaiman, S., Gibson, S., and Spencer, M., “Predictive dynamic digital holography and image sharpening,” *Journal of the Optical Society of America A* **35**(6), 923–935 (2018).
- [6] Ilhan, H. A., Dogar, M., and Ozcan, M., “Digital holographic microscopy and focusing methods based on image sharpness,” *Journal of Microscopy* **255**(3), 138–149 (2014).
- [7] Pellizari, C. J., Trahan III, R., Zhou, H., Williams, S., Williams, S. E., Nemati, B., Shao, M., and Bouman, C. A., “Synthetic aperture lidar: A model-based approach,” *IEEE Transactions on Computational Imaging* **3**(4), 901–916 (2017).
- [8] Bouman, C. A. and Sauer, K. D., “A unified approach to statistical tomography using coordinate descent optimization,” *IEEE Transactions on Image Processing* **5**(3), 480–492 (1996).
- [9] Bouman, C. A. and Sauer, K., “A generalized gaussian image model for edge-preserving map estimation,” *IEEE Transactions on Image Processing* **2**(3), 296–310 (1993).
- [10] Yu, Z., Thibault, J.-B., Sauer, K., Bouman, C. A., and Hsieh, J., “Accelerated line search for coordinate descent optimization,” *IEEE Nuclear Science Symposium* **5**, 2841–2844 (2006).
- [11] Schmidt, J., “Numerical simulation of optical wave propagation with examples in matlab,” *SPIE* (2010).
- [12] Srinath, S., Poyneer, L. A., Rudy, A. R., and Ammons, S. M., “Computationally efficient autoregressive method for generating phase screens with frozen flow and turbulence in optical simulations,” *Optical Society of America Optics Express* **23**(26) (2015).