

TIMBIR: A Method for Time-Space Reconstruction from Interlaced Views

K. Aditya Mohan, *Student Member, IEEE*, S. V. Venkatakrishnan, *Student Member, IEEE*, John W. Gibbs, E. Begum Gulsoy, Xianghui Xiao, Marc De Graef, Peter W. Voorhees, and Charles A. Bouman, *Fellow, IEEE*

Abstract—Synchrotron X-ray computed tomography (SXCT) is increasingly being used for 3D imaging of material samples at micron and finer scales. The success of these techniques has increased interest in 4D reconstruction methods that can image a sample in both space and time. However, the temporal resolution of widely used 4D reconstruction methods is severely limited by the need to acquire a very large number of views for each reconstructed 3D volume. Consequently, the temporal resolution of current methods is insufficient to observe important physical phenomena. Furthermore, measurement non-idealities also tend to introduce ring and streak artifacts into the 4D reconstructions.

In this paper, we present a time-interlaced model-based iterative reconstruction (TIMBIR) method which is a synergistic combination of two innovations. The first innovation, interlaced view sampling, is a novel method of data acquisition which distributes the view angles more evenly in time. The second innovation is a 4D model-based iterative reconstruction algorithm (MBIR) which can produce time-resolved volumetric reconstruction of the sample from the interlaced views. In addition to modeling both the sensor noise statistics and the 4D object, the MBIR algorithm also reduces ring and streak artifacts by more accurately modeling the measurement non-idealities. We present reconstructions of both simulated and real X-ray synchrotron data which indicate that TIMBIR can improve temporal resolution by an order of magnitude relative to existing approaches.

I. INTRODUCTION

Four-dimensional synchrotron X-ray computed tomography (4D-SXCT) is enabling scientists to study a wide variety of physical processes such as solidification and solid-state phase transformations in the field of material science [1], [2]. In contrast to conventional CT, 4D-SXCT produces time-resolved three-dimensional volumetric reconstruction of the sample.

Copyright (c) 2015 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

K. A. Mohan and C. A. Bouman are with the School of Electrical and Computer Engineering, Purdue University, 465 Northwestern Ave., West Lafayette, IN 47907-2035, USA. E-mail: {mohank,bouman}@purdue.edu.

E. B. Gulsoy and P. W. Voorhees are with the Department of Material Science and Engineering, Northwestern University, Evanston, IL, USA. E-mail: {e-gulsoy,p-voorhees}@northwestern.edu.

J. W. Gibbs is with the Los Alamos National Laboratory, Los Alamos, NM, USA. E-mail: jwgibbs@lanl.gov.

X. Xiao is with the Advanced Photon Source of Argonne National Laboratory, Lemont, IL, USA. E-mail: xhxiao@aps.anl.gov.

M. De Graef is with the Department of Material Science and Engineering, Carnegie Mellon University, Pittsburgh, PA, USA. E-mail: degraef@cmu.edu.

K. A. Mohan, S. Venkatakrishnan, E. B. Gulsoy, M. De Graef, P. W. Voorhees and C. A. Bouman were supported by an AFOSR/MURI grant #FA9550-12-1-0458. J. W. Gibbs was supported by a DOE NNSA Stewardship Science Graduate Fellowship (Grant No. DE-FC52-08NA28752). Use of the Advanced Photon Source, an Office of Science User Facility operated for the U.S. Department of Energy (DOE) Office of Science by Argonne National Laboratory, was supported by the U.S. DOE under Contract No. DE-AC02-06CH11357. This research was supported in part through computational resources provided by Information Technology at Purdue, West Lafayette, Indiana.

The high intensity and strong collimation of synchrotron radiation makes it especially suitable for high speed imaging of a wide variety of samples at the micron scale [1]–[3]. However, in-situ 4D imaging using SXCT still remains a major challenge owing to limitations on the data acquisition speed [1]. Moreover, impurities in the scintillator and imperfections in the detector elements cause ring artifacts in the reconstruction [4], [5]. Furthermore, detector pixels occasionally get saturated by high energy photons (often called “zingers”) which cause streak artifacts in the reconstruction [5], [6].




Fig. 1. Illustration of data acquisition for 4D-SXCT. A parallel beam of X-rays from a synchrotron is used to repeatedly image an object rotating at a constant speed. The intensity of the attenuated X-ray beam after passing through the object is measured by a 2D detector.

The traditional approach to 4D-SXCT is to acquire a sequence of parallel beam projections of the object, which is rotated at a constant speed, at progressively increasing equi-spaced view angles (henceforth called progressive view sampling) as shown in Fig. 1. Typically, the projections in each π radians rotation are grouped together and reconstructed into a single 3D volume using an analytic reconstruction algorithm such as filtered back projection (FBP) [7], [8] or a Fourier domain reconstruction method [9]–[11]. The time sequence of 3D reconstructions then forms the 4D reconstruction of the object.

Unfortunately, this traditional approach based on progressive view sampling and analytic reconstruction severely limits the temporal resolution of 4D reconstructions. The number of 3D volumes (henceforth called time samples) of the 4D reconstruction per unit time is given by $F_s = F_c/N_\theta$ where N_θ is the number of views used to reconstruct each time sample, and F_c is the data acquisition rate i.e., the rate at which projection images are collected. The maximum data acquisition rate, F_c , depends on a wide variety of hardware constraints such as the camera frame rate, the data transfer rate, buffer memory sizes, etc. The number of views, N_θ , required for spatial Nyquist sampling of the projection data is $\pi/2$ times the number of detector pixels, N_p , in the sensor’s field of view

perpendicular to the rotation axis [7]. However, in practice, for Nyquist sampling the number of views, N_θ , is chosen to be approximately equal to the number of cross-axial sampled pixels, N_p [7]. This means that in order to reconstruct a single time sample from a sensor with a frame size of 1024×1024 (i.e., a 1 mega-pixel sensor) one needs $N_\theta = 1024$ views of the object, resulting in a temporal reconstruction rate that is reduced by a factor of 1024 relative to the data rate, F_c . Consequently, the data rate becomes the limiting factor on the temporal resolution in a typical 4D-SXCT imaging system.

In order to increase the temporal reconstruction rate, either the number of projections per rotation can be reduced or the range of view angles used for reconstruction can be decreased. However, if the number of views per rotation is reduced the signal is under-sampled and analytic reconstruction algorithms produce substantial artifacts due to aliasing [7], [12]–[14]. If the number of view angles used for a single reconstruction is reduced, the Fourier space of the object is not fully sampled which results in a missing wedge of spatial frequencies [15]. With analytic reconstruction algorithms, this missing wedge results in time-varying non-isotropic resolution that will typically produce severe reconstruction artifacts. Therefore, using conventional reconstruction algorithms with a traditional progressive view sampling approach presents a fundamental limitation on the temporal resolution that can be achieved for a given spatial resolution.

In order to improve the quality of reconstruction, several new sampling strategies have been proposed for other tomographic applications. Using 2D Nyquist sampling theory, it has been shown that using a hexagonal sampling strategy [16]–[18] for the Radon transform the number of sample points can be reduced. Alternatively, by formulating the data acquisition as a time-sequential process, Willis et al. show that for an object with localized temporal variation the sampling rate can be reduced using an optimally scrambled angular sampling order [19], [20]. Another novel angular sampling strategy called equally sloped tomography [21], has been shown to produce superior quality reconstructions when reconstructing from a set of projections spaced equally in a slope parameter rather than angle. In [22], Zheng et al. present a method of reducing the number of projections by identifying favorable views based on prior knowledge of the object. In [23]–[27], the authors discuss different compressed sensing approaches for tomography. In general, the methods which change the angular order are optimal when the time interval between successive views is independent of their angular separation. Thus, optimal view sampling strategies have been shown to improve the reconstruction quality, even with conventional reconstruction algorithms.

An alternate approach to improving reconstruction quality is to use more advanced model-based iterative reconstruction (MBIR) methods, which are based on the estimation of a reconstruction which best fits models of both the sensor measurements (i.e., the forward model) and the object (i.e., prior model) [5], [28]–[30]. These 3D reconstruction methods have been shown to be very effective when the angular range is limited [31] and also when the number of views is less than that required by Nyquist sampling criterion [5]. In the

context of medical CT, several authors have also shown that modeling the temporal correlations [32]–[36] in addition to modeling the spatial correlations improves the quality of 4D MBIR reconstructions.

In this paper, we propose an approach to 4D reconstruction of time varying objects, which we call time interlaced model based iterative reconstruction (TIMBIR). TIMBIR is the synergistic combination of a novel interlaced view sampling technique with an innovative model-based iterative reconstruction (MBIR) algorithm. In [37] we present preliminary results using the TIMBIR method and in [38] we use TIMBIR to determine the morphology of a growing metallic dendrite in 4D. In the new interlaced view sampling method, all the views typically acquired over half a rotation using progressive view sampling are instead acquired over multiple half-rotations. Intuitively, interlaced view sampling spreads the view angles more evenly in time as opposed to the conventional progressive view sampling method that groups views at nearby angles together in time. Nonetheless, interlaced view sampling when used with conventional analytic reconstruction methods does not result in any gains since the number of views in each half-rotation is insufficient to achieve Nyquist sampling for a single reconstruction of the object. Consequently, analytic reconstruction methods produce severe artifacts when used to reconstruct the interlaced views at higher temporal rates.

In order to reconstruct the data acquired using the interlaced view sampling method, we propose a new 4D MBIR algorithm. In addition to modeling the measurement noise and spatio-temporal correlations in the 4D object, the MBIR algorithm reduces ring and streak artifacts by modeling the detector non-idealities [4], [6] and measurement outliers caused by high energy photons (called zingers) [5]. We adapt our forward model for 3D MBIR introduced in [5] to the current 4D framework and combine it with a modified q-generalized Gaussian Markov random field (qGGMRF) [39] based prior model, which models the spatio-temporal correlations in the reconstruction, and formulate the MBIR cost function. We note that our qGGMRF prior model is similar to the total-variation (TV) prior used in compressive sensing (CS) methods [24], [27]. The parameters of our qGGMRF prior model can be adjusted to result in the TV prior that is widely used in CS methods. We then present a fast distributed parallel multi-resolution algorithm based on surrogate functions to minimize this cost function.

The organization of this paper is as follows. In Section II, we present the interlaced view sampling technique. In Section III, we combine the synchrotron measurement model with the prior and formulate the MBIR cost function. In Section IV, we propose an algorithm to minimize the MBIR cost function. In Section V, we present a fast parallel distributed MBIR algorithm. In Section VI, we present results using both simulated and real datasets.

II. INTERLACED VIEW SAMPLING FOR TOMOGRAPHY

In order to satisfy the spatial Nyquist sampling requirement for each 3D time sample of a 4D reconstruction, it is typically necessary to collect approximately $N_\theta = N_p$ progressive

views, where N_p is the number of sampled pixels perpendicular to the axis of rotation [7]. In the traditional approach, these N_θ progressive views are taken in sequence while the sample is rotated continuously over π radians. The object is then reconstructed at a temporal rate of $F_s = F_c/N_\theta$ where F_c is the data acquisition rate.

In contrast to this approach, we propose an interlaced view sampling method where each frame of data consisting of N_θ distinct views are acquired over K interlaced sub-frames (see Fig. 2). Each sub-frame of data then consists of N_θ/K equally spaced views, but together the full frame of data contains all N_θ distinct views of the object. For a continuously rotating object, the formula which gives the view angle as a function of the discrete sample index, n , is given by

$$\theta_n = \left[nK + \mathcal{B}_r \left(\left\lfloor \frac{nK}{N_\theta} \right\rfloor \bmod K \right) \right] \frac{\pi}{N_\theta}, \quad (1)$$

where K is a power of 2, $b = \mathcal{B}_r(a)$ is the bit-reverse function which takes the binary representation of the integer a and reverses the order of the bits to form the output b [40]. It is interesting to note that there exist ordered subset methods [41] for 3D tomography which use a similar bit-reversal technique to determine the grouping of projection views. If we write a in base 2 expansion as $a = \sum_{i=0}^{l-1} b_i 2^i$ where $l = \log_2(K)$, then $b = \sum_{i=0}^{l-1} b_i 2^{l-1-i}$. If the angular range of the projections is limited to 2π radians, then the interlaced view angles are given by

$$\theta_n = \left[\left(n \bmod \frac{2N_\theta}{K} \right) K + \mathcal{B}_r \left(\left\lfloor \frac{nK}{N_\theta} \right\rfloor \bmod K \right) \right] \frac{\pi}{N_\theta}. \quad (2)$$

If the angular range of the projections is limited to π radians, then the interlaced view angles are given by

$$\theta_n = \left[\left(n \bmod \frac{N_\theta}{K} \right) K + \mathcal{B}_r \left(\left\lfloor \frac{nK}{N_\theta} \right\rfloor \bmod K \right) \right] \frac{\pi}{N_\theta}. \quad (3)$$

It is important to note that for the same index, n , the view angles generated by equations (1), (2), and (3) are all separated by an integer multiple of π radians. Hence, the projections obtained using equations (1), (2), and (3) are essentially the same. Fig. 2 compares progressive views with interlaced views and also highlights the interlacing of view angles across sub-frames. The object is then reconstructed at a temporal rate of $F_s = rF_c/N_\theta$ where r is the number of time samples of the 4D reconstruction in a frame. In TIMBIR, we reconstruct the object at a temporal rate which is r times the conventional rate of $F_s = F_c/N_\theta$. We typically set the parameter r equal to the number of sub-frames, K .

In TIMBIR, the conventional approach of reconstructing using progressive views is obtained when $r = K = 1$. We will show empirically that we can significantly improve the spatial and temporal reconstruction quality by increasing the value of K and r while using the same value of N_θ . Thus, we show that we can get significantly improved reconstruction quality by changing the view sampling method without increasing the amount of input data.

In progressive view sampling, the entire set of N_θ distinct views is acquired over a π radians rotation of the object. However, in interlaced view sampling all the distinct N_θ views

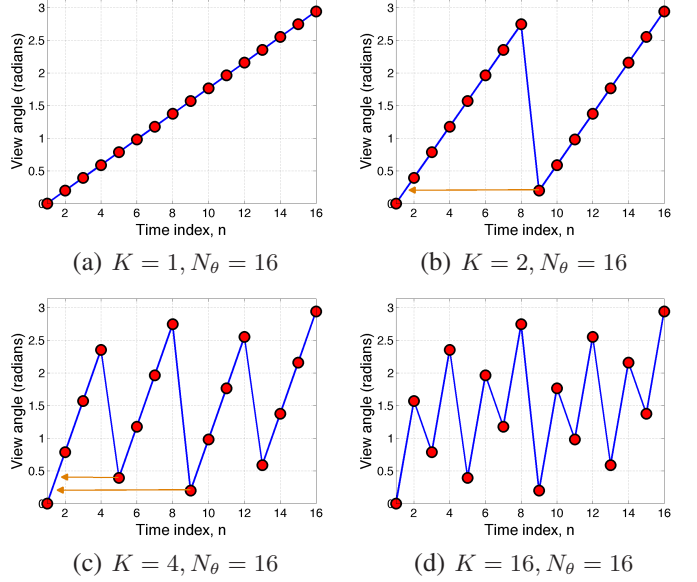


Fig. 2. Illustration of interlaced view sampling pattern for different values of K . (a-d) are plots of $\theta_n \bmod(\pi)$ vs. time index, n , for $K = 1, 2, 4$, and 16 . The arrows show the relative difference between the angular values across sub-frames.

are acquired over a $K\pi$ radians rotation of the object. Thus, if N_θ and F_c are fixed, then increasing the value of K increases the rotation speed of the sample. We will show that increasing the value of K can also improve the reconstruction quality. However, in some cases, the increased rotation speed of the sample may not be desirable. So in these cases, the parameter K can be adjusted to balance the need for improved image quality with the need to limit the sample's rotation speed. While increasing the rotation speed and reducing N_θ appears to be an intuitive step for increasing the temporal resolution of the reconstructions, we will empirically demonstrate that interlacing the views will give us significant improvements in the temporal reconstruction quality.

III. FORMULATION OF MBIR COST FUNCTION

The goal of SXCT is to reconstruct the attenuation coefficients of the sample from the acquired data. We reconstruct the attenuation coefficients of the object from the acquired data using the MBIR framework. The MBIR reconstruction is given by

$$(\hat{x}, \hat{\phi}) = \underset{x, \phi}{\operatorname{argmin}} \{ -\log p(y|x, \phi) - \log p(x) \}, \quad (4)$$

where $p(y|x, \phi)$ is the pdf of the projection data, y , given the object, x , and the unknown system parameters, ϕ , and $p(x)$ is a pdf for the 4D object.

A. Measurement Model

We begin by deriving a likelihood function $p(y|x, \phi)$ for the projections, y , from a time varying object, x . We model each voxel of the object as an independent piecewise constant function in time such that there are r equi-length reconstruction time samples in each frame. Thus, the projections ranging from $(j-1)N_\theta/r + 1$ to jN_θ/r are assumed to be generated from




Fig. 3. Plot of the generalized Huber function $\beta_{T,\delta}$ used in the likelihood term with $T = 3$ and $\delta = \frac{1}{2}$. Projections with large data mismatch error are penalized thereby reducing their influence in the overall cost function.

the j^{th} time sample. The vector of attenuation coefficients of the object at the j^{th} time sample is denoted by x_j .

A widely used model for X-ray transmission measurements is based on Beer's law and Poisson counting statistics for the measurement [42]. Using this model, if $\lambda_{n,i}$ is the measurement at the i^{th} detector element and n^{th} view and if $\lambda_{D,i}$ is the measurement in the absence of the sample, then an estimate of the projection integral is given by $y_{n,i} = \log\left(\frac{\lambda_{D,i}}{\lambda_{n,i}}\right)$. If we denote y to be the vector of projections $y_{n,i}$ and x to be the vector of attenuation coefficients at all time steps, then using a Taylor series approximation to the Poisson log-likelihood function [43] it can be shown that,

$$-\log p(y|x) \approx \frac{1}{2} \sum_{j=1}^L \sum_{n=n_{j-1}}^{n_j-1} \sum_{i=1}^M \left((y_{n,i} - A_{n,i,*} x_j) \frac{\sqrt{\Lambda_{n,i,i}}}{\sigma} \right)^2 + f(y), \quad (5)$$

where $n_j = j \frac{N_\theta}{r} + 1$, L is the total number of time samples in the reconstruction, $A_{n,i,*}$ is the i^{th} row of the forward projection matrix A_n , Λ_n is a diagonal matrix modeling the noise statistics, M is the total number of detector elements, and $f(y)$ is a constant that will be ignored in the subsequent optimization. The variance of the projection measurement, $y_{n,i}$, is inversely proportional to the mean photon count and hence we set $\Lambda_{n,i,i} = \lambda_{n,i}$ [42]. Since $\lambda_{n,i}$ is not equal to the photon count but is proportional to the photon count, there exists a constant of proportionality σ such that $\frac{\lambda_{n,i}}{\sigma^2}$ is the inverse variance of the projection measurement, $y_{n,i}$.

While this model is useful in several applications, it does not account for the non-idealities in the synchrotron measurement system. In particular, the log-likelihood term in (5) corresponds to a quadratic penalty on the weighted data mismatch error and does not account for the occurrence of zingers [5]. The zinger measurements correspond to a distribution with heavier tails than that corresponding to (5). Hence we change the quadratic penalty in (5) to a generalized Huber penalty (see Fig. 3) of the form [5], [44]

$$\beta_{T,\delta}(z) = \begin{cases} z^2 & |z| < T \\ 2\delta T|z| + T^2(1 - 2\delta) & |z| \geq T, \end{cases} \quad (6)$$

where T and δ are parameters of the generalized Huber function. In our model, the parameters of the generalized Huber function are chosen such that $0 < \delta < 1$ and $T > 0$. The generalized Huber function is non-convex in this range of parameter values. This penalty implies that if the ratio

of the data mismatch error to the noise standard deviation is greater than the parameter T , then the measured projection corresponds to a zinger.

Next, we model the non-idealities in the measurement system that cause ring artifacts. It has been shown [8] that the non-idealities that cause ring artifacts can be modeled via an additive detector dependent offset error, d_i , in the projection measurements, $y_{n,i}$. Hence, if we assume an unknown offset error d_i in the projection, $y_{n,i}$, then an estimate of the line integral is given by,

$$\tilde{y}_{n,i} = y_{n,i} - d_i. \quad (7)$$

The offset error, d_i , is typically not known from the measurements and hence we jointly estimate it during reconstruction. Thus, the new likelihood function that models the offset error and the zinger measurements is given by

$$p(y|x, d, \sigma) = \frac{1}{Z(\sigma)} \exp\{-U(y, x, d, \sigma)\}, \quad (8)$$

where $U(y, x, d, \sigma) =$

$$\frac{1}{2} \sum_{j=1}^L \sum_{n=n_{j-1}}^{n_j-1} \sum_{i=1}^M \beta_{T,\delta} \left((y_{n,i} - A_{n,i,*} x_j - d_i) \frac{\sqrt{\Lambda_{n,i,i}}}{\sigma} \right), \quad (9)$$

$Z(\sigma)$ is a normalizing constant and $d = [d_1 \cdots d_M]$ is the vector of all offset error parameters. Since (8) is a pdf, we can show that $Z(\sigma) = Z(1)\sigma^{ML\frac{N_\theta}{r}}$ using the property that (8) integrated over y should be equal to one [45]. Thus, the log-likelihood function is given by

$$-\log p(y|x, d, \sigma) = U(y, x, d, \sigma) + ML\frac{N_\theta}{r} \log(\sigma) + \tilde{f}(y), \quad (10)$$

where $\tilde{f}(y)$ is a constant which is ignored in the subsequent optimization. We note that when $\delta = 0$, $p(y|x, d, \sigma)$ is not a density function since it does not integrate to 1 and hence we assume $\delta > 0$ in the rest of the paper.

B. Prior Model

We use a q-generalized Gaussian Markov random field (qGGMRF) [39] based prior model for the voxels. The prior is used to model the 4D object in time as well as space. Using this model, the logarithm of the density function of x is given by

$$-\log p(x) = \sum_{j=1}^L \sum_{\{k,l\} \in \mathcal{N}} w_{kl} \rho_s(x_{j,k} - x_{j,l}) + \sum_{k=1}^P \sum_{\{j,i\} \in \mathcal{T}} \tilde{w}_{ji} \rho_t(x_{j,k} - x_{i,k}) + \text{constant}, \quad (11)$$

$$\text{where } \rho_s(\Delta) = \frac{\Delta_t \Delta_s^3 \left| \frac{\Delta}{\Delta_s \sigma_s} \right|^2}{c_s + \left| \frac{\Delta}{\Delta_s \sigma_s} \right|^{2-p}}, \quad \rho_t(\Delta) = \frac{\Delta_t \Delta_s^3 \left| \frac{\Delta}{\Delta_t \sigma_t} \right|^2}{c_t + \left| \frac{\Delta}{\Delta_t \sigma_t} \right|^{2-p}},$$

and $x_{j,k}$ is the k^{th} voxel of the object at time sample j , P is the total number of voxels in each time sample, \mathcal{N} is the set of all pairwise cliques in 3D space (all pairs of neighbors in a 26 point spatial neighborhood system), \mathcal{T} is the set of all pairs of indices of adjacent time samples (two

point temporal neighborhood system), p , c_s , c_t , σ_s and σ_t are qGGMRF parameters, Δ_s is a parameter proportional to the side length of a voxel and Δ_t is a parameter proportional to the duration of each time sample in the reconstruction. The weight parameters are set such that $w_{kl} \propto |k-l|^{-1}$, $\tilde{w}_{ji} \propto |j-i|^{-1}$, and normalized such that $\sum_{l \in \mathcal{N}_k} w_{kl} + \sum_{i \in \mathcal{T}_j} \tilde{w}_{ji} = 1$, where \mathcal{N}_k is the set of all spatial neighbors and \mathcal{T}_j is the set of all temporal neighbors of voxel $x_{j,k}$. The terms Δ_s and Δ_t in the prior model ensures invariance of the prior to changing voxel sizes [46].

C. MBIR Cost Function

By substituting (10) and (11) into (4), we get the following MBIR cost function,

$$c(x, d, \sigma) = \frac{1}{2} \sum_{j=1}^L \sum_{n=n_{j-1}}^{n_j-1} \sum_{i=1}^M \beta_{T,\delta} \left((y_{n,i} - A_{n,i,*} x_j - d_i) \frac{\sqrt{\Lambda_{n,i,i}}}{\sigma} \right) + \sum_{j=1}^L \sum_{\{k,l\} \in \mathcal{N}} w_{kl} \rho_s(x_{j,k} - x_{j,l}) + \sum_{k=1}^P \sum_{\{j,i\} \in \mathcal{T}} \tilde{w}_{ji} \rho_t(x_{j,k} - x_{i,k}) + ML \frac{N_\theta}{r} \log(\sigma). \quad (12)$$

The reconstruction is obtained by jointly minimizing the cost, $c(x, d, \sigma)$, with respect to x , d and σ . Additionally, we impose a linear constraint of the form $Hd = 0$ to minimize any shift in the mean value of the reconstruction. The form of the matrix H can be adjusted depending on the application. The form of the matrix H used in our application is described in appendix A.

IV. OPTIMIZATION ALGORITHM

The cost function (12) is non-convex in x , d and σ . Minimizing the current form of the cost function given by (12) is computationally expensive. So, instead we use the functional substitution approach [47], [48] to efficiently minimize (12). Our method also ensures monotonic decrease of the cost function (12). A substitute cost function $c_{\text{sub}}(x, d, \sigma; x', d', \sigma')$ to the cost function $c(x, d, \sigma)$ at the point (x', d', σ') is a function which bounds the cost function from above such that minimizing the substitute cost function results in a lower value of the original cost function.

A. Construction of Substitute Function

To derive a substitute function to the overall cost we find a substitute function to each term of the cost (12) and sum them together to derive an overall substitute function. In particular, we will use quadratic substitute functions, as they make the subsequent optimization computationally simple.

A sufficient condition for a function $q(z; z')$ to be a substitute function to $g(z)$ at the point z' is that $\forall z, q(z; z') \geq g(z)$ and $q(z'; z') = g(z')$. We can prove that

$$Q_{T,\delta}(z; z') = \begin{cases} z^2 & |z'| < T \\ \frac{\delta T}{|z'|} z^2 + \delta T |z'| + T^2(1 - 2\delta) & |z'| \geq T \end{cases}$$

```

function  $[x', e', b'] \leftarrow \text{UPDATEVOX}(x', e', b', \sigma', j, m)$ 
  for each voxel  $k$  in voxel-line  $m$  at time step  $j$  do
    Compute  $\tilde{\theta}_1$  and  $\tilde{\theta}_2$  as in equation (19)
     $\backslash\backslash$  Compute substitute function parameters
    for each  $l \in \mathcal{N}_k$  do
      Compute  $a_{jkl}^s$  using (14)
    end for
    for each  $i \in \mathcal{T}_j$  do
      Compute  $a_{kji}^t$  using (16)
    end for
    Compute  $\hat{x}_{j,k}$  as in equation (21)
     $u^* \leftarrow x'_{j,k} + \alpha(\hat{x}_{j,k} - x'_{j,k})$ 
     $\backslash\backslash$  Update all variables affected by update of  $x'_{j,k}$ 
     $e'_{j,n,i} \leftarrow e'_{j,n,i} - A_{n,i,j}(u^* - x'_{j,k}) \quad \forall n, i$ 
     $x'_{j,k} \leftarrow u^*$ 
    Using (18), update  $b'_{j,n,i} \quad \forall n, i$ 
  end for
end function

```

Fig. 4. Pseudo code to update a voxel. The voxel update is obtained by minimizing a quadratic substitute function in $x_{j,k}$ while the other variables are treated as constants.

```

function  $[\hat{x}, \hat{d}, \hat{\sigma}] \leftarrow \text{RECONSTRUCT}(y, x', d', \sigma')$ 
  %Inputs: Projections  $y$ , Initial reconstruction  $x'$ , Initial estimate of offset error  $d'$ , Initial estimate of variance parameter  $\sigma'$ 
  %Outputs: Reconstruction  $\hat{x}$ , estimate of offset error  $\hat{d}$ , estimate of variance parameter  $\hat{\sigma}$ 
   $e'_{j,n,i} = y_{n,i} - A_{n,i,*} x'_j - d'_i \quad \forall i, j, n$ 
  Set  $b'$  using (18)
  while Stopping criteria is not met do
     $\backslash\backslash$  Iteratively update all voxel lines in a random order
    for each voxel-line  $m$  at each time step  $j$  do
       $(x', e', b') \leftarrow \text{UPDATEVOX}(x', e', b', \sigma', j, m)$ 
    end for
     $d' \leftarrow$  Update  $d'$  using (24)
    Update  $\sigma'$  using (25)
  end while
   $\hat{x} \leftarrow x', \hat{d} \leftarrow d', \hat{\sigma} \leftarrow \sigma'$ 
end function

```

Fig. 5. Pseudo code of the 4D MBIR algorithm. The algorithm works by alternatively minimizing a substitute cost function (17) with respect to all the voxels, $x_{j,k}$, offset error, d_i , and variance parameter σ .

is a substitute function to $\beta_{T,\delta}(z)$ at the point z' by showing that it satisfies the sufficiency condition [44]. If the error sinogram is defined as $e_{j,n,i} = y_{n,i} - A_{n,i,*} x_j - d_i$ and $e'_{j,n,i} = y_{n,i} - A_{n,i,*} x'_j - d'_i$ is the error sinogram at the current values of (x', d', σ') , then a substitute function to $\beta_{T,\delta}(e_{j,n,i} \sqrt{\Lambda_{n,i,i}} / \sigma)$ in the original cost is given by $Q_{T,\delta}(e_{j,n,i} \sqrt{\Lambda_{n,i,i}} / \sigma; e'_{j,n,i} \sqrt{\Lambda_{n,i,i}} / \sigma')$ by Property 7.9 in [49].

A quadratic substitute function for the prior term $\rho_s(x_{j,k} - x_{j,l})$ can be shown to be [28],

$$\rho_s(x_{j,k} - x_{j,l}; x'_{j,k} - x'_{j,l}) = \frac{a_{jkl}^s}{2} (x_{j,k} - x_{j,l})^2 + b_{jkl}^s, \quad (13)$$

$$\text{where } a_{jkl}^s = \begin{cases} \frac{\rho'_s(x'_{j,k} - x'_{j,l})}{(x'_{j,k} - x'_{j,l})} & x'_{j,k} \neq x'_{j,l} \\ \rho''_s(0) & x'_{j,k} = x'_{j,l} \end{cases}, \quad (14)$$

and b_{jkl}^s is a constant. A quadratic substitute function for the prior term $\rho_t(x_{j,k} - x_{i,k})$ can be shown to be [28],

$$\rho_t(x_{j,k} - x_{i,k}; x'_{j,k} - x'_{i,k}) = \frac{a_{kji}^t}{2} (x_{j,k} - x_{i,k})^2 + b_{kji}^t, \quad (15)$$

$$\text{where } a_{kji}^t = \begin{cases} \frac{\rho'_t(x'_{j,k} - x'_{i,k})}{(x'_{j,k} - x'_{i,k})} & x'_{j,k} \neq x'_{i,k} \\ \rho''_t(0) & x'_{j,k} = x'_{i,k} \end{cases}, \quad (16)$$

and b_{kji}^t is a constant. Thus, a substitute function to (12) is given by,

$$\begin{aligned} c_{\text{sub}}(x, d, \sigma; x', d', \sigma') = & \frac{1}{2} \sum_{j=1}^L \sum_{n=n_{j-1}}^{n_j-1} \sum_{i=1}^M Q_{T,\delta} \left(e_{j,n,i} \frac{\sqrt{\Lambda_{n,i,i}}}{\sigma}; e'_{j,n,i} \frac{\sqrt{\Lambda_{n,i,i}}}{\sigma'} \right) \\ & + ML \frac{N_\theta}{r} \log(\sigma) + \sum_{j=1}^L \sum_{\{k,l\} \in \mathcal{N}} w_{kl} \rho_s(x_{j,k} - x_{j,l}; x'_{j,k} - x'_{j,l}) \\ & + \sum_{k=1}^P \sum_{\{j,i\} \in \mathcal{T}} \tilde{w}_{ji} \rho_t(x_{j,k} - x_{i,k}; x'_{j,k} - x'_{i,k}). \end{aligned} \quad (17)$$

B. Parameter Updates used in Optimization

To minimize the cost function given by (12), we repeatedly construct and minimize (17) with respect to each voxel, $x_{j,k}$, the offset error parameters, d_i , and the variance parameter, σ . To simplify the updates, we define $b'_{j,n,i}$ to be a indicator variable which classifies measurements as anomalous based on the current value of the error sinogram, $e'_{j,n,i}$, and the current value of the variance parameter, σ' , as shown below,

$$b'_{j,n,i} = \begin{cases} 1 & |e'_{j,n,i} \sqrt{\Lambda_{n,i,i}} / \sigma'| < T \\ 0 & |e'_{j,n,i} \sqrt{\Lambda_{n,i,i}} / \sigma'| \geq T \end{cases} \quad (18)$$

and let $\tilde{b}'_{j,n,i} = 1 - b'_{j,n,i}$.

1) *Voxel Update*: In order to minimize (17) with respect to a voxel k at time step j , we first rewrite (17) in terms of $x_{j,k}$ in the form of the following cost function,

$$\begin{aligned} \tilde{c}_{\text{sub}}(x_{j,k}) = & \tilde{\theta}_1 x_{j,k} + \frac{\tilde{\theta}_2}{2} (x_{j,k} - x'_{j,k})^2 \\ & + \sum_{l \in \mathcal{N}_k} w_{kl} \rho_s(x_{j,k} - x'_{j,l}; x'_{j,k} - x'_{j,l}) \\ & + \sum_{i \in \mathcal{T}_j} \tilde{w}_{ji} \rho_t(x_{j,k} - x'_{i,k}; x'_{j,k} - x'_{i,k}), \end{aligned} \quad (19)$$

where $\tilde{\theta}_1 = -\sum_{n=n_{j-1}}^{n_j-1} \sum_{i=1}^M A_{n,i,k} u'_{j,n,i}$, $\tilde{\theta}_2 = \sum_{n=n_{j-1}}^{n_j-1} \sum_{i=1}^M A_{n,i,k}^2 v'_{j,n,i}$,

$$\begin{aligned} u'_{j,n,i} = & \frac{\sqrt{\Lambda_{n,i,i}}}{\sigma'} \left[b'_{j,n,i} e'_{j,n,i} \frac{\sqrt{\Lambda_{n,i,i}}}{\sigma'} + \tilde{b}'_{j,n,i} \delta T \text{sgn}(e'_{j,n,i}) \right], \\ v'_{j,n,i} = & \frac{\sqrt{\Lambda_{n,i,i}}}{\sigma'} \left[b'_{j,n,i} \frac{\sqrt{\Lambda_{n,i,i}}}{\sigma'} + \tilde{b}'_{j,n,i} \frac{\delta T}{|e'_{j,n,i}|} \right], \end{aligned} \quad (20)$$

and sgn is the *signum* function. Then, the optimal update for $x_{j,k}$ is obtained by minimizing (19) with respect to $x_{j,k}$ and is given by

$$\hat{x}_{j,k} = \frac{\sum_{l \in \mathcal{N}_k} w_{kl} a_{jkl}^s x'_{j,l} + \sum_{i \in \mathcal{T}_j} \tilde{w}_{ji} a_{kji}^t x'_{i,k} + \tilde{\theta}_2 x'_{j,k} - \tilde{\theta}_1}{\sum_{l \in \mathcal{N}_k} w_{kl} a_{jkl}^s + \sum_{i \in \mathcal{T}_j} \tilde{w}_{ji} a_{kji}^t + \tilde{\theta}_2}. \quad (21)$$

In order to speed up the computation of voxel updates, we update certain groups of voxels sequentially. In particular, let w be the axis of rotation and $u - v$ axes be in the plane perpendicular to w . A voxel-line [28] consists of all voxels along w -axis which share the same $u - v$ coordinate. Since all voxels along a voxel-line share the same geometry computation in the $u - v$ plane, we update all the voxels along a voxel-line sequentially. To further speed up convergence of the algorithm, we use the over-relaxation method [28] which forces larger updates of voxels. The update for voxel $x_{j,k}$ is then given by

$$x'_{j,k} \leftarrow x'_{j,k} + \alpha (\hat{x}_{j,k} - x'_{j,k}), \quad (22)$$

where α is the over-relaxation factor which is set to be equal to 1.5. Typically, α is chosen in the range of (1, 2) which ensures decreasing values of (12) [28]. The pseudo code to update a voxel-line is shown in Fig. 4.

2) *Offset Error Update*: To minimize (17) with respect to d , subject to the constraint $Hd = 0$, we first rewrite (17) in terms of the offset error d as,

$$\tilde{c}_{\text{sub}}(d) = \frac{1}{2} \sum_{j=1}^L \sum_{n=n_{j-1}}^{n_j-1} (d - \bar{d}_{j,n})^t V'_{j,n} (d - \bar{d}_{j,n}) \quad (23)$$

where $\bar{d}_{j,n,i} = e'_{j,n,i} + d'_i$ and $V'_{j,n}$ is a diagonal matrix such that $V'_{j,n,i,i} = v'_{j,n,i}$ (given by (20)). We use the theory of Lagrange multipliers [50] to minimize (23) with respect to the offset error parameters, d_i , subject to the constraint $Hd = 0$. The optimal update for d is given by

$$\begin{aligned} d' \leftarrow & \Omega^{-1} \left(\sum_{j=1}^L \sum_{n=n_{j-1}}^{n_j-1} V'_{j,n} \bar{d}_{j,n} \right. \\ & \left. - H^t (H \Omega^{-1} H^t)^{-1} H \Omega^{-1} \sum_{j=1}^L \sum_{n=n_{j-1}}^{n_j-1} V'_{j,n} \bar{d}_{j,n} \right) \end{aligned} \quad (24)$$

where $\Omega = \sum_{j=1}^L \sum_{n=n_{j-1}}^{n_j-1} V'_{j,n}$.

3) *Variance Parameter Update*: The update for the variance parameter, σ^2 , is obtained by taking the derivative of (17) with respect to σ^2 and setting it to zero. The update is then given by,

$$\begin{aligned} \sigma'^2 \leftarrow & \frac{r}{N_\theta LM} \sum_{j=1}^L \sum_{n=n_{j-1}}^{n_j-1} \sum_{i=1}^M [e'^2_{j,n,i} \Lambda_{n,i,i} b'_{j,n,i} \\ & + \tilde{b}'_{j,n,i} \delta T |e'_{j,n,i}| \sigma' \sqrt{\Lambda_{n,i,i}}]. \end{aligned} \quad (25)$$

Thus, the pseudo code of the optimization algorithm to minimize (12) is shown in Fig. 5. The sequence of costs

is convergent since the surrogate function approach ensures monotonic decrease of the original cost function (12). We have also empirically observed that the reconstructions converge. In some cases, theoretical convergence proofs exist for majorization techniques with alternating minimization [51], [52]. However, due to the complicated nature of our cost function we have no theoretical proof of convergence for our algorithm.

Finally, we implemented non-homogeneous iterative coordinate descent (NHICD) [28] to improve the convergence speed of the algorithm. NHICD works by more frequently updating those voxels which have a greater need for updates. We also use multi-resolution initialization [44], [53] in which reconstructions at coarser resolutions are used to initialize reconstructions at finer resolution to improve convergence speed. The multi-resolution method transfers the computational load to the coarser scales where the optimization algorithm is faster due to the reduced dimensionality of the problem. Furthermore, we use bilinear interpolation to up-sample reconstructions at coarser resolutions to finer resolutions.

C. Algorithm Initialization

Since the MBIR cost function is non-convex, using a reasonable initial estimate is important to obtain a reasonable solution. At the coarsest scale of the multi-resolution algorithm, the variance parameter, σ , is initialized to one, the offset error, d_i , is set to zero and all voxels are initialized to zero. Furthermore, to prevent the algorithm from converging to a local minimum, we do not update the offset error, d_i , and variance parameter, σ , at the coarsest scale. At the i^{th} multi-resolution stage, the prior model parameter Δ_s is set to 2^{S-i} where S is the total number of multi-resolution stages and the parameter Δ_t is set to N_θ/r .

The algorithm stops when the ratio of average magnitude of voxel updates to the average magnitude of voxel values is less than a pre-defined threshold (convergence threshold). We use different values for the convergence threshold at different multi-resolution stages. If the convergence threshold at the finest multi-resolution stage is T , then the convergence threshold at the k^{th} multi-resolution stage is chosen to be $T/(S - k + 1)$.

V. DISTRIBUTED PARALLELIZATION OF MBIR

To enable high spatial and temporal resolution reconstructions typically required for 4D-SXCT, we propose a distributed parallel MBIR algorithm optimized to run on a high performance cluster (HPC) consisting of several supercomputing nodes. Each node consists of one or more multi-core processors with a single shared memory. In the distributed MBIR algorithm, the computation is distributed among several nodes so that the data dependency across nodes is minimized. If a node requires information from another node, then the information is explicitly communicated using a message passing interface (MPI). The multiple cores in each node are used to further speed up computation using OpenMP based shared memory parallelization.

Each iteration of the distributed MBIR algorithm consists of two update phases as shown in Fig. 6. All the voxels updated

in the first phase in Fig. 6 are not updated in the second phase and all the voxels updated in the second phase are not updated in the first phase. The algorithm alternates between the two update phases until convergence is achieved.

Due to the specific form of the voxel updates in (21), we need to ensure that only those voxels that do not share sinogram¹ entries or spatio-temporal neighbors are updated in parallel to ensure that the cost function decreases with each update. We equally divide all the slices in each time sample of the reconstruction (along w -axis) across multiple nodes as shown in Fig. 6. Each node only updates voxels in its share of the 3D volume at all the time samples. The projection data, y , is distributed such that each node only stores those projection slices in its local memory which depend on its share of the reconstruction. Such a distribution ensures that voxel updates in a node do not require sinogram entries from another node. Thus, there is no sinogram dependency across nodes. However, there is a voxel neighborhood dependency across nodes since the voxel slices which are on the “edge” of a 3D volume within a node (along w -axis) contain spatial neighbors which are stored in other nodes. Additionally, within a given node there are spatial and temporal neighbors which we need to ensure are not updated in parallel.

To address this problem, we first equally divide the slices in each node into an even number of blocks which we call “slice-blocks” as shown in Fig. 6. In each phase, the different darkly shaded slice-blocks in Fig. 6 are updated in parallel and the lightly shaded slice-blocks are updated in the next phase. The voxels within a slice-block are updated serially taking into advantage the same geometry computation along a “voxel-line” [28] as explained in the previous section. In phase 1, all the odd numbered slice-blocks at odd time samples and even numbered slice-blocks at even time samples are updated in parallel. In node k , after the updates of phase 1, the first slice of each odd time sample is communicated to node $k - 1$ and the last slice of each even time sample is communicated to node $k + 1$. In phase 2, all the even numbered slice-blocks at the odd time samples and the odd numbered slice-blocks at the even time samples are updated in parallel. In node k , after the updates of phase 2, the last slice of each odd time sample is communicated to node $k + 1$ and the first slice of each even time sample is communicated to node $k - 1$. All the slice-blocks within a node are updated in parallel using a OpenMP based shared memory parallelization scheme. This update scheme ensures that the cost-function (12) decreases monotonically with every iteration.

After updating the slice-blocks in each phase, only those elements of the error sinogram, e , and offset error, d , dependent on the last updated slice-blocks are updated. To update the variance parameter, σ , each node computes the summation in equation (25) corresponding to its share of the data and communicates the result to a master node. The master node then updates σ by accumulating the result from all the nodes and broadcasts the updated value back to the nodes. The updated values of slices from node $k - 1$ and node $k + 1$

¹A sinogram image is a column wise stack of a single projection slice (i.e., projection values along a single detector line) across all the views.

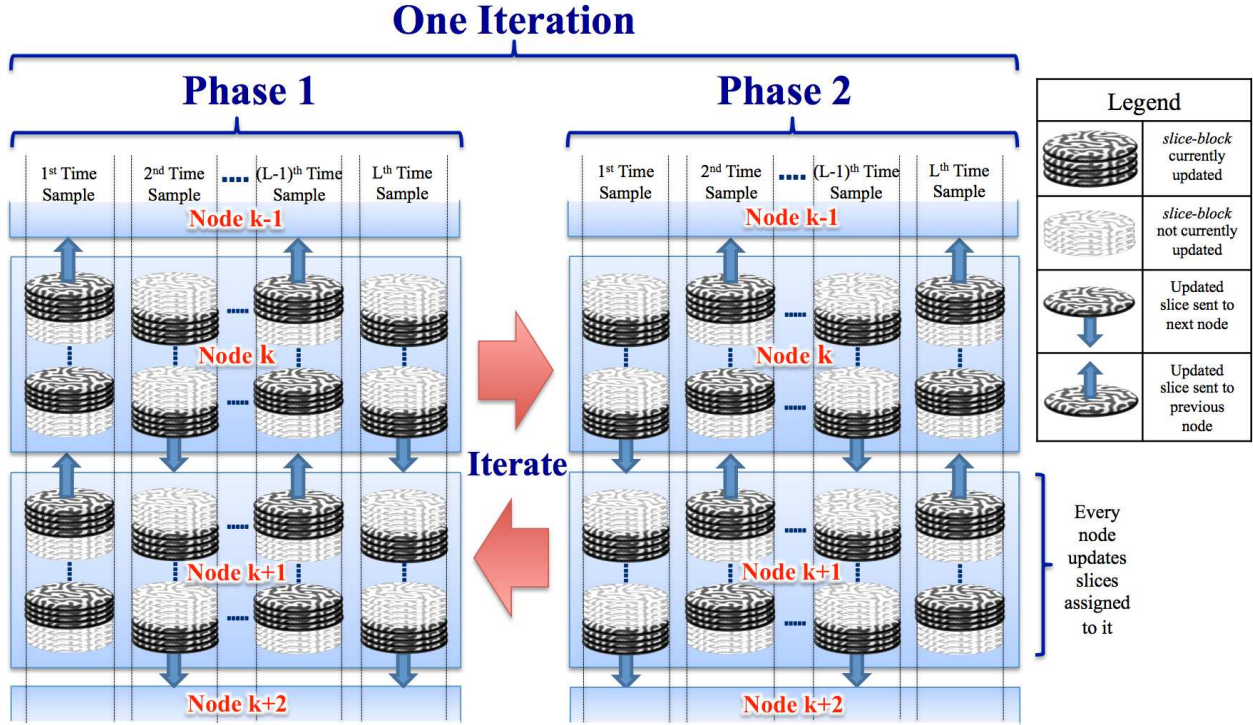


Fig. 6. Block diagram describing the distributed parallel MBIR algorithm. The data y is split axially among multiple nodes. Each node then updates only those slices of the reconstruction that depend on its share of the data. In phase 1, all the even numbered slice-blocks are updated at even time samples and all the odd numbered slice blocks are updated at odd time samples. In phase 2, all the odd numbered slice-blocks are updated at even time samples and all the even numbered slice blocks are updated at odd time samples. The algorithm iterates between the two phases until convergence is achieved.

required by node k are communicated using MPI calls. The offset error, d , is updated when the slices are communicated across nodes. Thus, the compute nodes are not idle during communication. In all our experiments, the time taken for computing the offset error update was found to be greater than the time taken for communicating the slices. Thus, the algorithm is limited by computational speed and not by the time taken for communication and synchronization.

VI. EXPERIMENTAL RESULTS

A. Simulated Data Set

In this section, we compare FBP and MBIR reconstructions of simulated datasets using both the traditional progressive view sampling and the proposed interlaced view sampling methods. The simulated dataset is generated from a time-varying phantom by accurately modeling the data acquisition in a real physical system. First, a 2D phantom is generated using the Cahn-Hilliard equation which models the process of phase separation in the cross-axial plane ($u-v$ axes) [54]. The two phases of the object have attenuation coefficients of 2.0 mm^{-1} and 0.67 mm^{-1} respectively. The 3D phantom is then generated by repeatedly stacking the 2D phantom along the axial dimension (w -axis). This phantom is representative of the phenomenon that we are interested in studying in 4D. A $u-v$ slice of the phantom at different time instants is shown in Fig. 7 (a - c) and Fig. 7 (d) shows a v -axis slice of the phantom (along the red line in Fig. 7 (a)) as a function of time. The phantom is assumed to have a voxel resolution of $0.65 \times 0.65 \times 0.65 \text{ } \mu\text{m}^3$ and a size of $N_w \times N_v \times N_u = 16 \times 1024 \times 1024$. The phantom is sampled in time at the

data acquisition rate, F_c , and the projections are generated by forward projecting the sampled phantom at the appropriate angles. To simulate the detector non-idealities in SXCT, we add an offset error d_i to the projection $y_{n,i}$ at every n^{th} view. To simulate the effect of zingers, we randomly set 0.1% of the projections, y , to zero. We also simulate the measurement data $\lambda_{n,i}$ to have Poisson statistics. The simulated value of the variance parameter is $\sigma^2 = 10$.

The simulated sensor has a resolution of $N_p = 256$ pixels in the cross-axial direction and 4 pixels in the axial direction. A 3D time sample of the 4D reconstruction has a voxel resolution of $2.6 \times 2.6 \times 2.6 \text{ } \mu\text{m}^3$ and a size of $N_w \times N_v \times N_u = 4 \times 256 \times 256$. The temporal reconstruction rate is $F_s = rF_c/N_\theta$ where r is the number of time samples of the reconstruction in a frame, F_c is the data acquisition rate, and N_θ is the number of distinct views acquired. Since the temporal reconstruction rate varies with r , the reconstructions are up sampled in time to the data acquisition rate, F_c , using cubic interpolation and then compared with the phantom. Also, since the phantom has higher spatial resolution than the reconstructions, the phantom is down-sampled by averaging over blocks of pixels to the reconstruction resolution before comparison. The regularization parameters in the prior model are chosen such that they minimize the root mean square error (RMSE) between the reconstruction and the phantom. The parameter p of the qGGMRF model is set to 1.2 and the parameters of the generalized Huber function are set to be $\delta = 0.5$ and $T = 4$. We use a convergence threshold of $T = 0.01$.

The traditional approach to 4D-SXCT is to use progressive




Fig. 7. The conventional approach to reconstruction using progressive views. (a-d) shows the time-varying phantom with two attenuation phases. Reconstruction of progressive views with $N_\theta = 256$ using FBP is shown in (e-h) and using MBIR is shown in (i-l) ($K = 1, N_\theta = 256, r = 1$). The first three columns show a $u-v$ slice of the object sampled at different time instants. The last column shows the time evolution of a v -axis slice of the object (along the red line in (a)). Reconstruction of rapidly time-varying objects using conventional methods results in poor temporal resolution and spatial artifacts.

view sampling and an analytic reconstruction algorithm such as FBP. So, we first generate a dataset of progressive views satisfying the Nyquist spatial sampling criterion and reconstruct it using FBP and MBIR algorithms. For analytic algorithms, the Nyquist view sampling requirement for a sensor of size $N_p = 256$ in the cross-axial dimension is to acquire projections at $N_\theta = 256$ distinct angles using progressive view sampling. Thus, we generate a dataset of progressive views with $N_\theta = 256$ and then reconstruct it at a temporal rate of $F_s = F_c/256$ by reconstructing $r = 1$ time sample every frame. The FBP reconstruction of this dataset is shown in Fig. 7 (e - h) and the MBIR reconstruction is shown in Fig. 7 (i - l). When compared to FBP, MBIR produces lower noise reconstructions while preserving the spatial resolution (Fig. 7 (e - g) and Fig. 7 (i - k)). Furthermore, FBP reconstructions suffer from strong ring and streak artifacts. From Fig. 7 (h) & (l), we can also see that neither FBP nor MBIR are able to reconstruct temporal edges accurately with progressive views.

Next, we investigate different methods of increasing temporal resolution using progressive views. First, we reconstruct a dataset of progressive views satisfying the Nyquist criterion at a temporal rate faster than the conventional method. Thus, we reconstruct a progressive view dataset with parameter $N_\theta = 256$ using MBIR at a rate of $F_s = F_c/32$ by reconstructing $r = 8$ time samples every frame. As shown in Fig. 8 (a - c), the reconstruction obtained using this method

has strong artifacts due to a missing wedge of view angles in the data used to reconstruct every time sample of the reconstruction. In the next approach, we reduce the number of distinct angles in the progressive view dataset to $N_\theta = 32$ and reconstruct at a rate of $F_s = F_c/32$ by reconstructing $r = 1$ time sample every frame. In this case, due to the severe under sampling of views, the MBIR reconstruction suffers from severe loss in quality as shown in Fig. 8 (e - h). This illustrates that merely reducing the number of views every π radians and using an advanced reconstruction algorithm such as MBIR is insufficient for our problem. Thus, there is an inherent sub-optimality in using progressive view sampling to achieve high temporal resolution reconstructions.

Finally, we investigate the effect of interlaced views on the spatial and temporal reconstruction quality when reconstructing using FBP and MBIR algorithms. We reconstruct a dataset of interlaced views in which each frame of $N_\theta = 256$ distinct angles is interlaced over $K = 8$ sub-frames. The object is then reconstructed at a rate of $F_s = F_c/32$ by reconstructing $r = 8$ time samples every frame. The reconstruction using FBP is shown in Fig. 8 (i - l) and using MBIR is shown in Fig. 8 (m - p). We can see that reconstructing the interlaced views with FBP results in extremely poor quality reconstructions. In contrast, MBIR with interlaced views (TIMBIR) results in a substantially better reconstruction of the object with minimal artifacts as shown in Fig. 8 (m - p). Furthermore, we can see

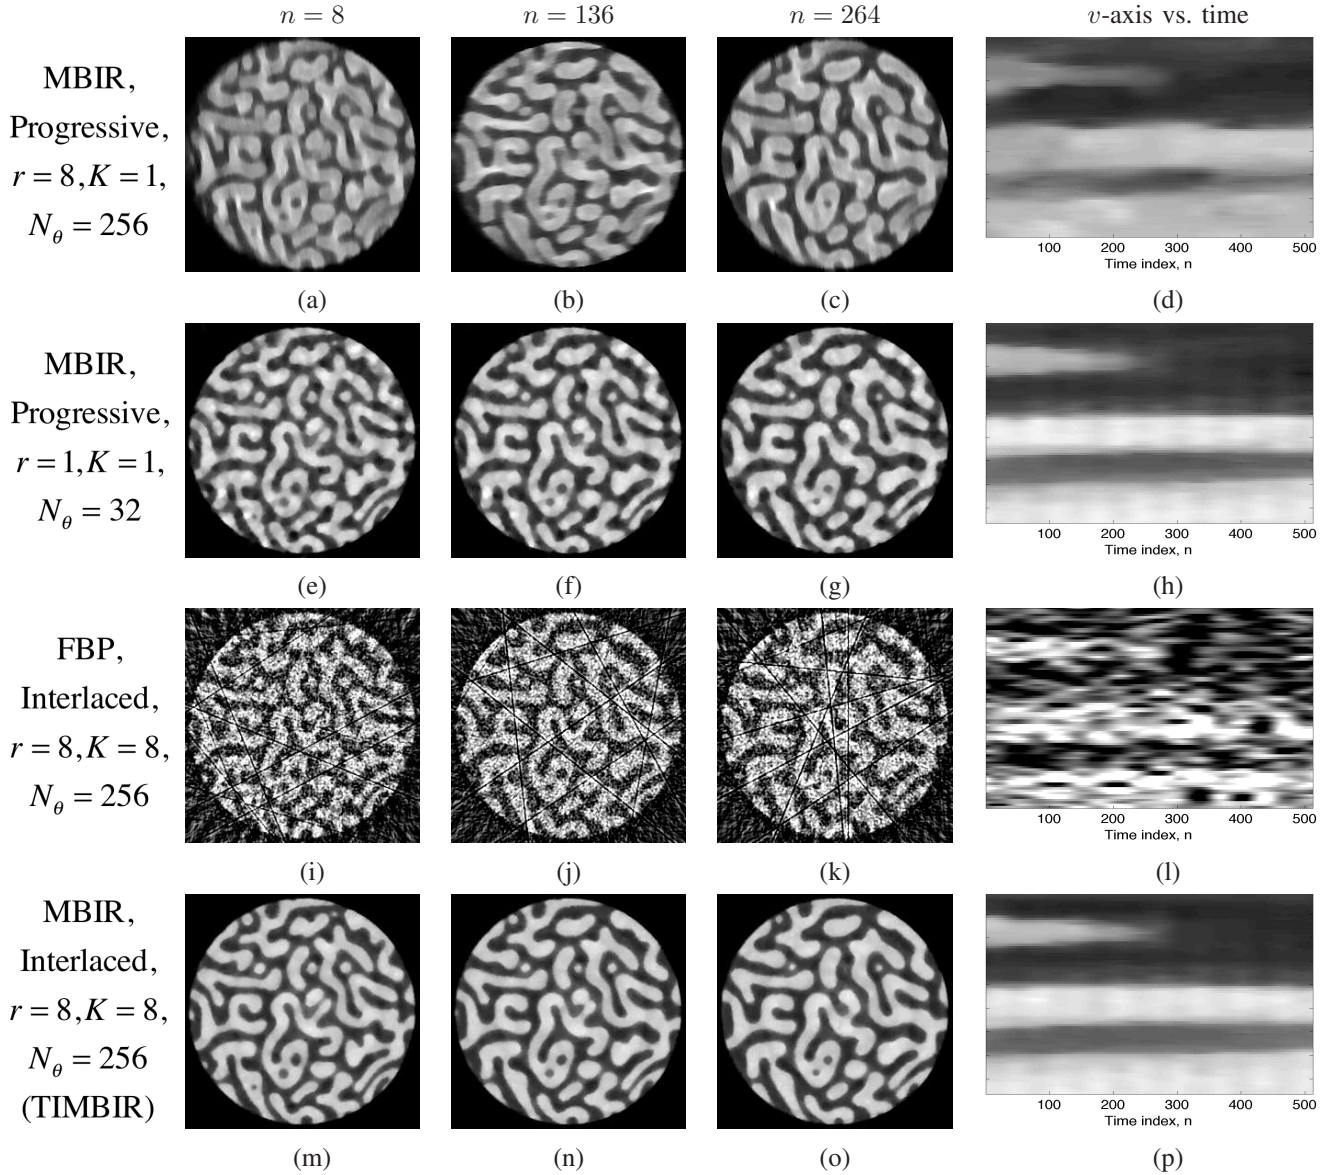


Fig. 8. Comparison of TIMBIR with other approaches to high temporal resolution reconstruction. The first three columns show a $u-v$ slice of the 4D reconstruction at different times and the last column shows a v -axis slice of the reconstruction versus time. All reconstructions have a temporal reconstruction rate of $F_s = F_c/32$. (a-d) is MBIR reconstruction ($r = 8$) of progressive views with $N_\theta = 256$. (e-h) is MBIR reconstruction ($r = 1$) of progressive views with $N_\theta = 32$. The reconstruction ($r = 8$) of interlaced views with $K = 8, N_\theta = 256$ using FBP is shown in (i-l) and using MBIR (TIMBIR) is shown in (m-p). Progressive view sampling results in poor reconstruction quality in time and/or space. Moreover, interlaced view sampling combined with FBP still causes severe artifacts due to under-sampling of view angles. However, TIMBIR produces high quality reconstructions in both time and space.




Fig. 9. Illustration of the effect of K on the RMSE between the reconstruction and the phantom when $r = 8$ and $N_\theta = 256$. Since the RMSE reduces as K is increased, interlacing of views is vital when reconstructing at a higher temporal rate of $F_s = F_c/32$ ($r = 8$).

that TIMBIR is able to more accurately reconstruct temporal

edges (Fig. 8 (p)) than other methods (Fig. 8 (d, h & l) and Fig. 7 (h,l)). Thus, by comparing Fig. 8 (m - p) with Fig. 8 (a - h) and Fig. 7 (i - l), we can conclude that interlaced view sampling is superior to progressive view sampling when reconstructing using MBIR. Furthermore, by comparing Fig. 8 (i - l) with Fig. 7 (e - h), we can conclude that combining interlaced view sampling with FBP does not result in any improvements. In Fig. 10, we plot a single voxel as a function of time for different values of r, K , and N_θ . We can see that TIMBIR produces the most accurate reconstruction of the voxel as a function of time among all the methods. The root mean squared errors (RMSE) between the reconstructions and the phantom ground-truth shown in Table I support these visual conclusions. Thus, TIMBIR with its synergistic combination of interlaced sampling and MBIR reconstruction results in a

TABLE I

ROOT MEAN SQUARE ERROR BETWEEN THE RECONSTRUCTION AND THE PHANTOM. TIMBIR HAS THE LOWEST RMSE AMONG ALL THE METHODS

Description	Parameters	RMSE (mm ⁻¹)
Nyquist Progressive, FBP ($F_s = F_c/256$)	$K = 1, N_\theta = 256$ $r = 1$	0.3376
Nyquist Progressive, MBIR ($F_s = F_c/256$)	$K = 1, N_\theta = 256$ $r = 1$	0.2182
Nyquist Progressive, MBIR ($F_s = F_c/32$)	$K = 1, N_\theta = 256$ $r = 8$	0.3302
Sub-Nyquist Progressive, MBIR ($F_s = F_c/32$)	$K = 1, N_\theta = 32$ $r = 1$	0.2393
Interlaced, FBP ($F_s = F_c/32$)	$K = 8, N_\theta = 256$ $r = 8$	0.6035
Interlaced, MBIR ($F_s = F_c/32$) (TIMBIR)	$K = 8, N_\theta = 256$ $r = 8$	0.2119

much higher quality reconstruction than either method can achieve by itself.




Fig. 10. Plot of a voxel as a function of time for different values of r , K , and N_θ . The most accurate reconstruction of the voxel as a function of time is obtained for the case of TIMBIR with parameters $r = 8$, $K = 8$, and $N_\theta = 256$.

Fig. 9 shows the RMSE between the MBIR reconstructions and the phantom as a function of K when $N_\theta = 256$ and $r = 8$. We can see that the RMSE reduces as the number of sub-frames, K , is increased from 1 to 8. This shows that interlacing of views is vital to achieving high temporal resolution reconstructions.

In the supplementary document, we have presented additional results comparing the segmentation quality of the reconstructions using different methods. We have also presented reconstructions with ring and streak artifacts obtained by using the standard forward model given in (5).

B. Real Data Set

To demonstrate the performance gains achieved by TIMBIR in a real physical system, we reconstruct the dendritic growth in an Al-Cu alloy [1] in 4D. Dendrites are complex tree like structures which form as liquids are cooled from a sufficiently high temperature. It is of great interest to study dendritic growth since the morphology of the growing dendrites determine the properties of many materials. Thus, by studying

the dendritic growth we can better understand the processes controlling the morphology of these tree like structures.

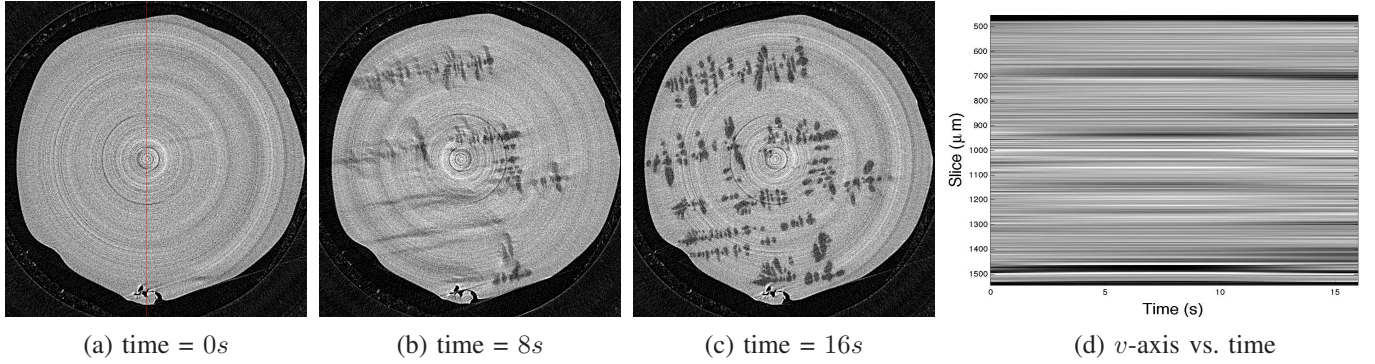
In our experiments, the data acquisition rate is limited to $F_c = 192$ Hz due to limitations on the camera frame rate, the data transfer rate, and buffer sizes. The detector width is $N_p = 1600$ pixels in the cross-axial direction and 1080 pixels along the axial direction with a pixel resolution of $0.65 \mu\text{m} \times 0.65 \mu\text{m}$. However, we only reconstruct a window of 4 pixels in the axial direction and $N_p = 1536$ pixels in the cross-axial direction. The reconstructions have a size of $N_w \times N_v \times N_u = 4 \times 1536 \times 1536$ and a voxel resolution of $0.65 \times 0.65 \times 0.65 \mu\text{m}^3$. The exposure time of the detector is set to 1 ms. The regularization parameters, σ_s and σ_t , are chosen to provide the best visual reconstruction quality. The parameter p of the qGGMRF model is set to 1.2 and the parameters of the generalized Huber function are set to be $\delta = 0.5$ and $T = 4$. We use a convergence threshold of $T = 0.05$. We increased the threshold from $T = 0.01$ used for simulated data to $T = 0.05$ for real data to reduce run time.

Imaging is done using polychromatic X-ray radiation from a synchrotron. The attenuated X-rays from the object are converted to visible light using a $25 \mu\text{m}$ thick LuAG:Ce scintillator which is then imaged by a PCO Edge CMOS camera. We also use a $10\times$ magnifying objective to get an effective pixel resolution of $0.65 \mu\text{m} \times 0.65 \mu\text{m}$. The polychromatic radiation results in beam hardening which causes cupping artifacts in the reconstruction. To correct for these artifacts, we use a simple quadratic correction polynomial of the form $y = az^2 + z$ where z is the measured value of the projection and y is the projection after correction [55]. The corrected projections are then used to do the reconstruction. To find the optimal value of a , we observed the reconstruction for cupping artifacts for increasing values of a . Based on this empirical analysis, we used a value of $a = 0.5$ which was found to minimize cupping artifacts.

First, we reconstruct the 4D object using a progressive view dataset satisfying the Nyquist criterion. To acquire N_θ progressive views over a rotation of π radians at a data acquisition rate of F_c , the object has to be rotated by an angle of π radians every $F_c^{-1} N_\theta$ seconds. Thus, the angular rotation speed of the object is fixed at $R_s = \pi F_c / N_\theta = 0.3925$ radians per second and data is acquired at a rate of $F_c = 192$ Hz using progressive view sampling with parameter $N_\theta = 1536$. The object is then reconstructed at a rate of 0.125 Hz by reconstructing $r = 1$ time sample every frame. With this method, we get a temporal resolution of $F_s^{-1} = 8s$. The FBP and MBIR reconstructions of data acquired using this technique are shown in Fig. 11 (a - d) and Fig. 11 (e - h) respectively. From Fig. 11 (d, h), we can see that a temporal resolution of $8s$ is inadequate to temporally resolve the growing dendrites. Furthermore, it also causes blur artifacts in the reconstructions as seen in Fig. 11 (b, f). The strong ring artifacts in the FBP reconstructions results in additional distortion. However, by modeling the measurement non-idealities, MBIR is able to substantially reduce ring artifacts. Thus, using the conventional method we cannot reconstruct the dendritic growth with sufficient resolution in time and space.

To increase the temporal resolution with progressive views,

FBP, Progressive Views ($r = 1, K = 1, N_\theta = 1536$)



MBIR, Progressive Views ($r = 1, K = 1, N_\theta = 1536$)

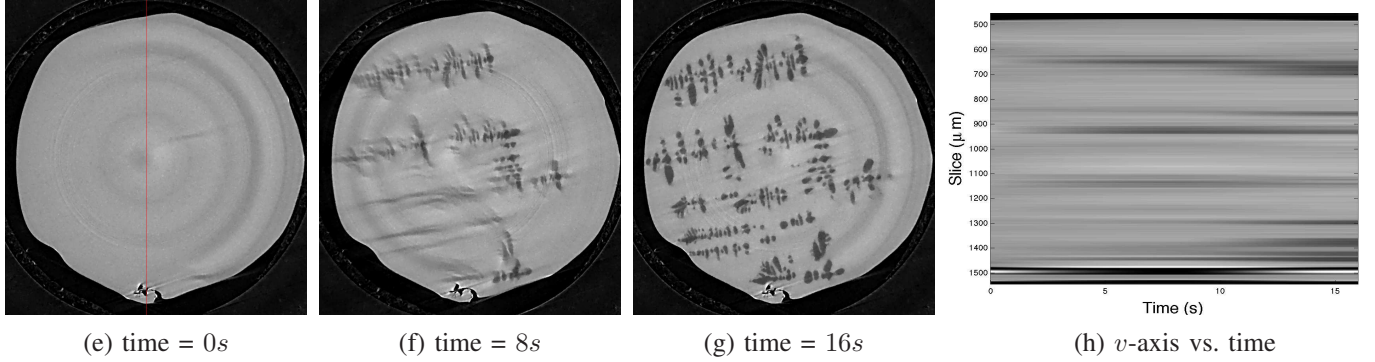



Fig. 11. Reconstructions of dendritic growth using Nyquist progressive views. The first three columns show a $u - v$ slice of the sample at different times. The last column shows a v -axis slice of the sample (along the red line in (a, e)) as a function of time. The reconstruction ($r = 1$) from progressive views with $N_\theta = 1536$ (Nyquist criteria) using FBP is shown in (a - d) and using MBIR is shown in (e - h). The reconstruction shown in (a - h) has a low temporal reconstruction rate of $F_s = 0.125$ Hz which is insufficient to temporally resolve the growing dendrites.

MBIR, Progressive Views ($r = 1, K = 1, N_\theta = 48$)



TIMBIR : MBIR, Interlaced Views ($r = 32, K = 32, N_\theta = 1536$)




Fig. 12. Comparison of reconstructions of dendritic growth using TIMBIR with other approaches to high temporal resolution reconstructions. The first three columns show a $u - v$ slice of the sample at different times. The last column shows a v -axis slice of the sample (along the red line in (a, e)) as a function of time. The MBIR reconstruction ($r = 1$) from progressive views with $N_\theta = 48$ (Sub-Nyquist) is shown in (a - d). The MBIR reconstruction ($r = 8$) from interlaced views with $K = 32, N_\theta = 1536$ is shown in (e - h) (TIMBIR). Even though the reconstruction in (a - d) has a rate of $F_s = 4$ Hz, the spatial reconstruction quality is very poor. However, the TIMBIR reconstruction in (e - h) not only has a rate of $F_s = 4$ Hz to temporally resolve the growing dendrites, it also has good spatial reconstruction quality. Note that results in (a-d) and (e-h) correspond to data from different experiments.

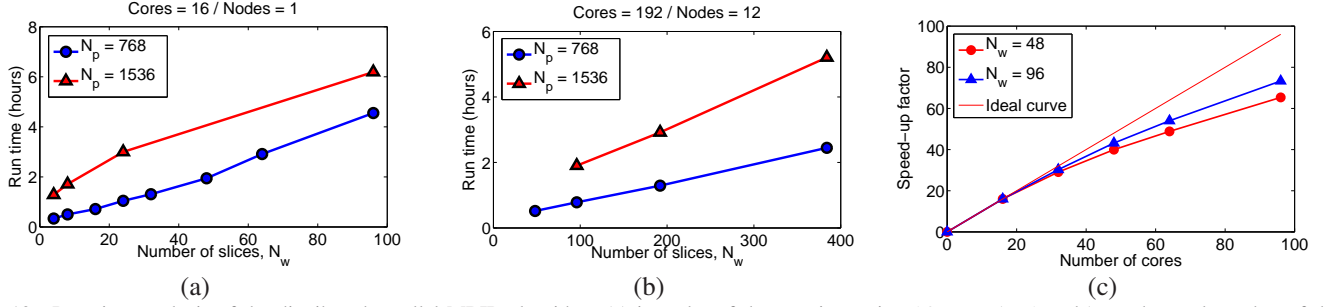


Fig. 13. Run time analysis of the distributed parallel MBIR algorithm. (a) is a plot of the run time using 16 cores (or 1 node) vs. the total number of the slices, N_w , at different cross-axial detector resolutions, N_p . (b) is a plot of the run time using 192 cores (or 12 nodes) vs. the total number of slices, N_w . (c) is a plot of the run time speed-up factor vs. the number of cores. The speed-up factor with n cores is computed as the ratio of the average run time per iteration with 1 core to that with n cores. We can see that the algorithm speed up factor gradually deviates from the ideal behavior when the number of cores is increased. Furthermore, the speed up behavior improves when the number of slices, N_w , is increased.

we reduce the number of distinct angles to $N_\theta = 48$ (Sub-Nyquist) per π radians, increase the rotation speed to $R_s = 12.56$ radians per second, and acquire data at the maximum rate of $F_c = 192$ Hz. The object is then reconstructed at a rate of $F_s = 4$ Hz by reconstructing $r = 1$ time sample every frame. The reconstructions using this method are shown in Fig. 12 (a - d). Note that the dendritic structure is different than in Fig. 11 (a - h) because it is not possible to replicate such physical phenomenon in an exact manner with each experimental run. With this method, we get a temporal resolution of $F_s^{-1} = 0.25$ s. In this case, even though the temporal resolution is high (Fig. 12 (d)), the spatial reconstruction quality is very poor (Fig. 12 (a - c)).

Next, we study the effect of TIMBIR on the reconstruction quality when reconstructing at a higher temporal rate of $F_s = 4$ Hz. To acquire N_θ interlaced views over a rotation of $K\pi$ radians at a rate of F_c , the object has to be rotated by an angle of π radians every $F_c^{-1}N_\theta/K$ seconds. Thus, the angular rotation speed of the object is fixed at $R_s = \pi K F_c / N_\theta = 12.56$ radians per second and data is acquired at a rate of $F_c = 192$ Hz using interlaced view sampling with parameters $K = 32$ and $N_\theta = 1536$. The object is then reconstructed at a rate of $F_s = 4$ Hz by reconstructing $r = 32$ time samples every frame. The reconstructions using this method are shown in Fig. 12 (e - h). We can see that the reconstruction using TIMBIR has very good spatial (Fig. 12 (e - g)) and temporal resolution (Fig. 12 (h)). Thus, TIMBIR with its synergistic combination of interlaced view sampling and MBIR algorithm is able to reconstruct the dendritic evolution at a high spatial and temporal resolution. We have empirically found that using a value of K less than but closest to $\sqrt{N_\theta}$ gives us good results.

C. Computational Cost

We study the variation of run time of the MBIR algorithm with the number of cores when reconstructing one frame of projections with parameters $K = 32$, $N_\theta = 1536$, $r = 32$. The algorithm run time is determined for two different cross-axial detector resolutions of $N_p = 768$ and 1536 and for different number of axial slices, N_w . The cross-axial reconstruction resolution is $N_v \times N_u = N_p \times N_p$ and the number of slices in the reconstruction is N_w . To determine the run time, we run the algorithm on the Conte supercomputing cluster at Purdue

University. Each node on Conte consists of two 8 core Intel Xeon-E5 processors. Fig. 13 (a, b) shows the run time as a function of the number of slices, N_w , for two different cross-axial detector resolutions of $N_p = 768$ and 1536 . The run time using 16 cores (or 1 node) is shown in Fig. 13 (a) and using 192 cores (or 12 nodes) is shown in Fig. 13 (b). Thus, we can see that increasing the number of cores from 16 to 192 significantly reduces the run time of the MBIR algorithm.

In Fig. 13 (c), we plot the speed-up factor as a function of the number of cores for different number of slices when $N_p = 768$. The speed-up factor with n cores is computed as the ratio of the average run time per iteration with 1 core to that with n cores at the finest multi-resolution stage. From the figure, we can see that initially the speed-up improvement is almost linear with the number of cores and gradually deviates from the ideal behavior when the number of cores is increased. Furthermore, when the number of slices, N_w , is increased, the speed-up behavior improves since the computation can be more efficiently distributed among the different nodes. Thus, using the distributed MBIR algorithm we can efficiently reconstruct large volumes using multiple cores. In the future, we believe the run time can be significantly reduced using better parallel architectures.

VII. CONCLUSION

In this paper, we propose a novel interlaced view sampling method which when combined with our 4D MBIR algorithm is able to achieve a synergistic improvement in reconstruction quality of 4D-SXCT. In addition to accounting for spatial and temporal correlations in the object, the MBIR algorithm also accounts for the measurement non-idealities encountered in a SXCT system. Using the new interlaced view sampling strategy with the 4D MBIR algorithm (TIMBIR), we were able to achieve a $32\times$ improvement in temporal resolution by reconstructing $32\times$ more time samples while preserving spatial reconstruction quality. We also present a distributed parallel MBIR algorithm to enable reconstructions of large datasets. In the future, by using better spatial and temporal prior models we expect to achieve much better reconstruction quality than presented in this paper. Furthermore, the TIMBIR method can be extended to any tomographic technique involving sampling of projections at different angles.

APPENDIX A

THE OFFSET ERROR CONSTRAINT MATRIX, H

In this section, we briefly discuss our choice of the H matrix. We choose the matrix H such that we enforce a zero constraint on the weighted average of the offset errors, d_i , over overlapping rectangular patches. Since the number of constraints is much less than the number of parameters, d_i , the matrix H will have many more columns than rows. Furthermore, all the rows of H sum to the same value and all the columns sum to the same value.

The offset error constraint $Hd = 0$ is better expressed as a constraint on the two dimensional form of the offset error, $\tilde{d}_{i,j}$. Let $\tilde{d}_{i,j}$ be the offset error corresponding to the i^{th} row and j^{th} column of the detector. We then impose a zero constraint on the weighted average of the offset error, $\tilde{d}_{i,j}$, over overlapping² rectangular patches as shown in Fig. 14. The patches are such that they overlap half-way along both the u -axis and v -axis and cover the entire detector plane.




Fig. 14. Diagram showing the overlapping rectangular patches over which the weighted sum of the offset error is zero. The patches are such that they overlap half-way along both the u -axis and v -axis and cover the entire detector plane.

Let $\tilde{h}_N(i)$ be a triangular window of the form

$$\tilde{h}_N(i) = \begin{cases} i & 1 \leq i \leq N \\ 2N - i + 1 & N + 1 \leq i \leq 2N \\ 0 & \text{otherwise.} \end{cases} \quad (26)$$

Then, the function $h(i, j) = \tilde{h}_P(i)\tilde{h}_Q(j)$ is used to appropriately weight the offset error terms, $d_{i,j}$, over the rectangular patches. Let $h^{(k,l)}(i, j)$ be the weighting function for the k^{th} patch along the u -axis and l^{th} patch along the v -axis obtained by appropriately shifting $h(i, j)$. Then, the constraint corresponding to the (k, l) patch is given by,

$$\sum_{i=0}^{M_r} \sum_{j=0}^{M_c} h^{(k,l)}(i, j) \tilde{d}_{i,j} = 0. \quad (27)$$

where M_r and M_c are the total number of rows and columns of the detector respectively. In our application, we choose P such that it is closest to $\sqrt{M_r}$ and is a factor of M_r (similarly, we choose Q depending on M_c). Thus, the number of constraints

is approximately equal to $\sqrt{M_r M_c}$. Each constraint indexed by (k, l) in (27) corresponds to one of the rows of H . The offset error d is the vector form of $\tilde{d}_{i,j}$ and the elements of H correspond to the weights, $h^{(k,l)}(i, j)$, used in (27). Thus, the matrix H has a size of $\sqrt{M_r M_c} \times M_r M_c$.

REFERENCES

- [1] D. Rowenhorst and P. Voorhees, "Measurement of interfacial evolution in three dimensions," *Annual Review of Materials Research*, vol. 42, no. 1, pp. 105–124, 2012.
- [2] L. K. Aagesen, J. L. Fife, E. M. Lauridsen, and P. W. Voorhees, "The evolution of interfacial morphology during coarsening: A comparison between 4D experiments and phase-field simulations," *Scripta Materialia*, vol. 64, no. 5, pp. 394–397, 2011.
- [3] H. A. Bale, A. Haboub, A. A. MacDowell, J. R. Nasiatka, D. Y. Parkinson, B. N. Cox, D. B. Marshall, and R. O. Ritchie, "Real-time quantitative imaging of failure events in materials under load at temperatures above 1,600° C," *Nature Materials*, vol. 12, no. 1, 2013.
- [4] J. Sijbers and A. Postnov, "Reduction of ring artifacts in high resolution micro-CT reconstructions," *Physics in Medicine and Biology*, vol. 49, no. 14, p. N247, 2004.
- [5] K. A. Mohan, S. V. Venkatakrishnan, L. F. Drummy, J. Simmons, D. Y. Parkinson, and C. A. Bouman, "Model-based iterative reconstruction for synchrotron X-ray tomography," in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*, May 2014, pp. 6909–6913.
- [6] T. Donath, "Quantitative X-ray microtomography with synchrotron radiation," Ph.D. dissertation, GKSS-Forschungszentrum, 2007.
- [7] A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. Society for Industrial and Applied Mathematics, 2001.
- [8] J. Hsieh, *Computed Tomography: Principles, Design, Artifacts, and Recent Advances*. SPIE Press, Nov. 2009.
- [9] D. Gottlieb, B. Gustafsson, and P. Forsén, "On the direct Fourier method for computer tomography," *Medical Imaging, IEEE Transactions on*, vol. 19, no. 3, pp. 223–232, 2000.
- [10] F. Marone and M. Stampanoni, "Regridding reconstruction algorithm for real-time tomographic imaging," *Journal of Synchrotron Radiation*, vol. 19, no. 6, pp. 1029–1037, Nov 2012.
- [11] J. Walden, "Analysis of the direct Fourier method for computer tomography," *Medical Imaging, IEEE Transactions on*, vol. 19, no. 3, pp. 211–222, 2000.
- [12] P. J. La Riviere and X. Pan, "Noise properties of periodic interpolation methods with implications for few-view tomography," *Nuclear Science, IEEE Transactions on*, vol. 46, no. 3, pp. 639–645, 1999.
- [13] M. Bertram, J. Wiegert, D. Schafer, T. Aach, and G. Rose, "Directional view interpolation for compensation of sparse angular sampling in cone-beam CT," *Medical Imaging, IEEE Transactions on*, vol. 28, no. 7, pp. 1011–1022, 2009.
- [14] R. R. Galigekere, K. Wiesent, and D. W. Holdsworth, "Techniques to alleviate the effects of view aliasing artifacts in computed tomography," *Medical Physics*, vol. 26, no. 6, 1999.
- [15] M. Davison, "The ill-conditioned nature of the limited angle tomography problem," *SIAM Journal on Applied Mathematics*, vol. 43, no. 2, pp. 428–448, 1983.
- [16] P. Rattey and A. G. Lindgren, "Sampling the 2-D Radon transform," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 29, no. 5, pp. 994–1002, 1981.
- [17] A. Faridani and L. R. Erik, "High-resolution computed tomography from efficient sampling," *Inverse Problems*, vol. 16, no. 3, p. 635, 2000.
- [18] A. M. Cormack, "Sampling the Radon transform with beams of finite width," *Physics in Medicine and Biology*, vol. 23, no. 6, p. 1141, 1978.
- [19] N. Willis and Y. Bresler, "Optimal scan for time-varying tomography I: Theoretical analysis and fundamental limitations," *Image Processing, IEEE Transactions on*, vol. 4, no. 5, pp. 642–653, 1995.
- [20] —, "Optimal scan for time-varying tomography II: Efficient design and experimental validation," *Image Processing, IEEE Transactions on*, vol. 4, no. 5, pp. 654–666, 1995.
- [21] J. Miao, F. Förster, and O. Levi, "Equally sloped tomography with oversampling reconstruction," *Physical Review B*, vol. 72, no. 5, p. 052103, 2005.
- [22] Z. Zheng and K. Mueller, "Identifying sets of favorable projections for few-view low-dose cone-beam CT scanning," in *11th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine*, 2011.

²If the patch extends outside the support of the detector plane along any dimension, then we wrap it around from the start of that dimension.

- [23] Y. Kaganovsky, D. Li, A. Holmgren, H. Jeon, K. P. MacCabe, D. G. Polite, J. A. O'Sullivan, L. Carin, and D. J. Brady, "Compressed sampling strategies for tomography," *J. Opt. Soc. Am. A*, vol. 31, no. 7, pp. 1369–1394, Jul 2014.
- [24] H. Yu and G. Wang, "Compressed sensing based interior tomography," *Physics in Medicine and Biology*, vol. 54, no. 9, p. 2791, 2009.
- [25] S. Leng, J. Tang, J. Zambelli, B. Nett, R. Tolakanahalli, and G.-H. Chen, "High temporal resolution and streak-free four-dimensional cone-beam computed tomography," *Physics in Medicine and Biology*, vol. 53, no. 20, p. 5653, 2008.
- [26] G.-H. Chen, P. Theriault-Lauzier, J. Tang, B. Nett, S. Leng, J. Zambelli, Z. Qi, N. Bevin, A. Raval, S. Reeder, and H. Rowley, "Time-resolved interventional cardiac C-arm cone-beam CT: An application of the PICCS algorithm," *Medical Imaging, IEEE Transactions on*, vol. 31, no. 4, pp. 907–923, April 2012.
- [27] R. Leary, Z. Saghi, P. A. Midgley, and D. J. Holland, "Compressed sensing electron tomography," *Ultramicroscopy*, vol. 131, no. 0, pp. 70–91, 2013.
- [28] Z. Yu, J. Thibault, C. Bouman, K. Sauer, and J. Hsieh, "Fast model-based X-ray CT reconstruction using spatially nonhomogeneous ICD optimization," *IEEE Trans. on Image Processing*, vol. 20, no. 1, pp. 161–175, Jan. 2011.
- [29] S. V. Venkatakrishnan, L. F. Drummy, M. A. Jackson, M. De Graef, J. Simmons, and C. A. Bouman, "A model based iterative reconstruction algorithm for high angle annular dark field-scanning transmission electron microscope (HAADF-STEM) tomography," *Image Processing, IEEE Transactions on*, vol. 22, no. 11, pp. 4532–4544, 2013.
- [30] Z. Ruoqiao, J. B. Thibault, C. A. Bouman, K. D. Sauer, and H. Jiang, "Model-based iterative reconstruction for dual-energy X-ray CT using a joint quadratic likelihood model," *Medical Imaging, IEEE Transactions on*, vol. 33, no. 1, pp. 117–134, 2014.
- [31] S. J. Kisner, E. Haneda, C. A. Bouman, S. Skatter, M. Kourinny, and S. Bedford, "Limited view angle iterative CT reconstruction," *Proc. SPIE*, vol. 8296, pp. 82 960F–82 960F–9, 2012.
- [32] H. Wu, A. Maier, R. Fahrig, and J. Horneegger, "Spatial-temporal total variation regularization (STTVR) for 4D-CT reconstruction," in *Proceedings of the SPIE*, vol. 8313, pp. 83 133J–83 133J–7, 10.1117/12.911162.
- [33] L. Ritschl, S. Sawall, M. Knaup, A. Hess, and M. Kachelriess, "Iterative 4D cardiac micro-CT image reconstruction using an adaptive spatio-temporal sparsity prior," *Physics in Medicine and Biology*, vol. 57, no. 6, p. 1517, 2012.
- [34] X. Jia, Y. Lou, B. Dong, Z. Tian, and S. Jiang, "4D computed tomography reconstruction from few-projection data via temporal non-local regularization," in *Proceedings of the 13th International Conference on Medical Image Computing and Computer-assisted Intervention: Part I*, ser. MICCAI'10. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 143–150.
- [35] Z. Tian, X. Jia, B. Dong, Y. Lou, and S. B. Jiang, "Low-dose 4DCT reconstruction via temporal nonlocal means," *Medical Physics*, vol. 38, no. 3, pp. 1359–1365, 2011.
- [36] J. Hinkle, M. Szegedi, B. Wang, B. Salter, and S. Joshi, "4D CT image reconstruction with diffeomorphic motion model," *Medical Image Analysis*, vol. 16, no. 6, pp. 1307–1316, 2012.
- [37] K. A. Mohan, S. V. Venkatakrishnan, J. W. Gibbs, E. B. Gulsoy, X. Xiao, M. De Graef, P. W. Voorhees, and C. A. Bouman, "4D model-based iterative reconstruction from interlaced views," *accepted for publication at IEEE International Conference on Acoustics Speech and Signal Processing (ICASSP)*, April 2015.
- [38] J. W. Gibbs, K. A. Mohan, E. B. Gulsoy, A. J. Shahani, X. Xaio, C. Bouman, M. D. Graef, and P. W. Voorhees, "The three-dimensional morphology of growing dendrites," *submitted to Nature Scientific Reports*.
- [39] J.-B. Thibault, K. D. Sauer, C. A. Bouman, and J. Hsieh, "A three-dimensional statistical approach to improved image quality for multislice helical CT," *Medical Physics*, vol. 34, no. 11, pp. 4526–4544, 2007.
- [40] R. M. Cramblitt and J. P. Allebach, "Analysis of time-sequential sampling with a spatially hexagonal lattice," *J. Opt. Soc. Am.*, vol. 73, no. 11, pp. 1510–1517, Nov 1983.
- [41] D. Kim, S. Ramani, and J. Fessler, "Combining ordered subsets and momentum for accelerated X-ray CT image reconstruction," *Medical Imaging, IEEE Transactions on*, vol. 34, no. 1, pp. 167–178, Jan 2015.
- [42] K. Sauer and C. Bouman, "Bayesian Estimation of Transmission Tomograms Using Segmentation Based Optimization," *IEEE Trans. on Nuclear Science*, vol. 39, pp. 1144–1152, 1992.
- [43] C. Bouman and K. Sauer, "A unified approach to statistical tomography using coordinate descent optimization," *IEEE Trans. on Image Processing*, vol. 5, no. 3, pp. 480–492, Mar. 1996.
- [44] S. Venkatakrishnan, L. Drummy, M. Jackson, M. De Graef, J. Simmons, and C. Bouman, "Model-based iterative reconstruction for bright field electron tomography," *Computational Imaging, IEEE Transactions on*, no. 99, pp. 1–1, 2014.
- [45] S. Saquib, C. Bouman, and K. Sauer, "ML parameter estimation for Markov random fields with applications to Bayesian tomography," *Image Processing, IEEE Transactions on*, vol. 7, no. 7, pp. 1029–1044, Jul 1998.
- [46] S. Oh, A. Milstein, C. Bouman, and K. Webb, "A general framework for nonlinear multigrid inversion," *Image Processing, IEEE Transactions on*, vol. 14, no. 1, pp. 125–140, 2005.
- [47] J. Zheng, S. Saquib, K. Sauer, and C. Bouman, "Parallelizable Bayesian tomography algorithms with rapid, guaranteed convergence," *Image Processing, IEEE Transactions on*, vol. 9, no. 10, pp. 1745–1759, 2000.
- [48] H. Erdogan and J. Fessler, "Monotonic algorithms for transmission tomography," *Medical Imaging, IEEE Transactions on*, vol. 18, no. 9, pp. 801–814, 1999.
- [49] C. A. Bouman, *Model Based Image And Signal Processing*, 2013.
- [50] E. K. Chong and S. H. Zak, *An introduction to optimization*. John Wiley & Sons, 2013, vol. 76.
- [51] M. Jacobson and J. Fessler, "An expanded theoretical treatment of iteration-dependent majorize-minimize algorithms," *Image Processing, IEEE Transactions on*, vol. 16, no. 10, pp. 2411–2422, Oct 2007.
- [52] M. Razaviyayn, M. Hong, and Z. Luo, "A unified convergence analysis of block successive minimization methods for nonsmooth optimization," *SIAM Journal on Optimization*, vol. 23, no. 2, pp. 1126–1153, 2013.
- [53] M. Kamasak, C. Bouman, E. Morris, and K. Sauer, "Direct reconstruction of kinetic parameter images from dynamic PET data," *IEEE Trans. on Medical Imaging*, vol. 24, no. 5, pp. 636–650, May 2005.
- [54] J. W. Cahn and J. E. Hilliard, "Free energy of a nonuniform system. I. interfacial free energy," *The Journal of Chemical Physics*, vol. 28, no. 2, 1958.
- [55] P. Jin, C. Bouman, and K. Sauer, "A method for simultaneous image reconstruction and beam hardening correction," in *IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, Oct 2013, pp. 1–5.



K. Aditya Mohan (S'08) is currently a PhD candidate at the department of electrical and computer engineering, Purdue University, IN, USA. He received his B.Tech. degree in electronics and communication engineering from National Institute of Technology Karnataka, Surathkal, India, in 2010 and M.S. degree in electrical and computer engineering from Purdue University, West Lafayette, IN, in 2014.

From 2010 to 2011, he worked as an ASIC design engineer at Nvidia, Bengaluru, India. His research interests are in statistical signal processing, inverse problems, computational imaging, computed tomography, and computer vision.



S. V. Venkatakrishnan (S'12) received his B.Tech. degree in electronics and communication engineering from National Institute of Technology Tiruchirappalli, India, in 2007, M.S. and Ph.D. degrees in electrical and computer engineering from Purdue University, West Lafayette, IN, in 2009, and 2014 respectively.

From 2009 to 2010, he worked as a research and development engineer at Baker Hughes Inc., Houston, TX on logging while drilling - imaging magnetometers. His research interests are in statistical information processing, inverse problems, computational imaging and machine learning.




John W. Gibbs is currently a director's funded postdoc at Los Alamos National Laboratory, Los Alamos, NM, USA. He received BS and MS degrees in metallurgical engineering from Colorado School of Mines in 2008 and 2009, and a PhD from Northwestern University in materials science and engineering in 2014. His research interests include phase transformations in metals, pattern formation during solidification, and in-situ characterization of solidification via x-ray radiography and tomography.



Emine B. Gulsoy is the Associate Director of the Center for Hierarchical Materials Design, a consortium lead by Northwestern University, University of Chicago and Argonne National Laboratory. She received the B.S. degree from Sabanci University, Istanbul, Turkey, in 2006. She received the M.S. and Ph.D. degrees from Department of Materials Science and Engineering, Carnegie Mellon University, Pittsburgh, PA, USA, in 2008 and 2010 respectively, with her thesis titled Computational Tools for Automated Analysis of Three-dimensional

Microstructural Characterization Data. She joined Department of Materials Science and Engineering at Northwestern University, Evanston, IL, USA, as a Post-Doctoral Research Fellow in 2010, where she currently is a Research Associate. She also holds a joint-affiliation with the X-Ray Sciences Division of Argonne National Laboratory, Argonne, IL, USA. Her research interests include multi-modal imaging and microstructural characterization using serial sectioning via optical and electron microscopy and X-ray tomography, with an emphasis on coarsening in metal alloys and energy storage materials.



Xianghui Xiao received the B.S. degree in metallurgical engineering from the Chongqing University, Chongqing, China, in 1993, and the Ph.D. degree in physics from the Institute of High Energy Physics, Chinese Academy of Sciences, China, in 2002. He was a Post-Doctoral Researcher with the Cornell High Energy Synchrotron Source, Ithaca, NY, USA, and Advanced Photon Source of Argonne National Laboratory (APS/ANL) before he became a beam-line scientist at APS/ANL. His current research interests include applying time-resolving tomography

techniques in diverse dynamic phenomena studies, and develop suitable hardware and software in such applications.



Marc De Graef received his BS and MS degrees in physics from the University of Antwerp (Belgium) in 1983, and his Ph.D. in physics from the Catholic University of Leuven (Belgium) in 1989, with a thesis on copper-based shape memory alloys. He then spent three and a half years as a post-doctoral researcher in the Materials Department at the University of California at Santa Barbara before joining Carnegie Mellon in 1993. He is currently professor and co-director of the J. Earle and Mary Roberts Materials Characterization Laboratory. His research

areas are in materials characterization by means of electron microscopy and X-ray tomography techniques. Prof. De Graef is a Fellow of the Microscopy Society of America, and received the 2012 Educator Award from The Minerals, Metals, and Materials Society.



Peter Voorhees is the Frank C. Engelhart Professor of Materials Science and Engineering at Northwestern University, and Professor of Engineering Sciences and Applied Mathematics. He is co-director of the Northwestern-Argonne Institute of Science and Engineering and is co-director of the Center for Hierarchical Materials Design. He received his Ph.D. in Materials Engineering from Rensselaer Polytechnic Institute. He was a member of the Metallurgy Division at the National Institute for Standards and Technology until joining the Department of Materials Science and Engineering at Northwestern University in 1988. He has received numerous awards including the National Science Foundation Presidential Young Investigator Award, ASM International Materials Science Division Research Award (Silver Medal), the TMS Bruce Chalmers Award, the ASM J. Willard Gibbs Phase Equilibria Award, the McCormick School of Engineering and Applied Science Award for Teaching Excellence, and is listed as a Highly Cited Researcher by the Institute for Scientific Information. Professor Voorhees is a fellow of ASM International, the Minerals, Metals and Materials Society, and the American Physical Society. He has published over 200 papers in the area of the thermodynamics and kinetics of phase transformations.



Charles A. Bouman (S'86-M'89-SM'97-F'01) received a B.S.E.E. degree from the University of Pennsylvania in 1981 and a MS degree from the University of California at Berkeley in 1982. From 1982 to 1985, he was a full staff member at MIT Lincoln Laboratory and in 1989 he received a Ph.D. in electrical engineering from Princeton University. He joined the faculty of Purdue University in 1989 where he is currently the Michael J. and Katherine R. Birck Professor of Electrical and Computer Engineering. He also holds a courtesy appointment in the School of Biomedical Engineering and is co-director of Purdue's Magnetic Resonance Imaging Facility located in Purdue's Research Park.

Professor Bouman's research focuses on the use of statistical image models, multiscale techniques, and fast algorithms in applications including tomographic reconstruction, medical imaging, and document rendering and acquisition. Professor Bouman is a Fellow of the IEEE, a Fellow of the American Institute for Medical and Biological Engineering (AIMBE), a Fellow of the society for Imaging Science and Technology (IS&T), a Fellow of the SPIE professional society. He is also a recipient of IS&T's Raymond C. Bowman Award for outstanding contributions to digital imaging education and research, has been a Purdue University Faculty Scholar, and received the College of Engineering Engagement/Service Award, and Team Award. He was previously the Editor-in-Chief for the IEEE Transactions on Image Processing and a Distinguished Lecturer for the IEEE Signal Processing Society, and he is currently a member of the Board of Governors. He has been an associate editor for the IEEE Transactions On Image Processing and the IEEE Transactions On Pattern Analysis and Machine Intelligence. He has also been Co-Chair of the 2006 SPIE/IS&T Symposium on Electronic Imaging, Co-Chair of the SPIE/IS&T conferences on Visual Communications and Image Processing 2000 (VCIP), a Vice President of Publications and a member of the Board of Directors for the IS&T Society, and he is the founder and Co-Chair of the SPIE/IS&T conference on Computational Imaging.