Accelerating image reconstruction in three-dimensional optoacoustic tomography on graphics processing units

Kun Wang and Chao Huang

Department of Biomedical Engineering, Washington University in St. Louis, St. Louis, Missouri 63130

Yu-Jiun Kao and Cheng-Ying Chou

Department of Bio-Industrial Mechatronics Engineering, National Taiwan University, Taipei 106, Taiwan

Alexander A. Oraevsky

TomoWave Laboratories, 675 Bering drive, Suite 575, Houston, Texas, Houston, Texas 77057

Mark A. Anastasio^{a)}

Department of Biomedical Engineering, Washington University in St. Louis, St. Louis, Missouri 63130

(Received 19 September 2012; revised 19 November 2012; accepted for publication 20 December 2012; published 28 January 2013)

Purpose: Optoacoustic tomography (OAT) is inherently a three-dimensional (3D) inverse problem. However, most studies of OAT image reconstruction still employ two-dimensional imaging models. One important reason is because 3D image reconstruction is computationally burdensome. The aim of this work is to accelerate existing image reconstruction algorithms for 3D OAT by use of parallel programming techniques.

Methods: Parallelization strategies are proposed to accelerate a filtered backprojection (FBP) algorithm and two different pairs of projection/backprojection operations that correspond to two different numerical imaging models. The algorithms are designed to fully exploit the parallel computing power of graphics processing units (GPUs). In order to evaluate the parallelization strategies for the projection/backprojection pairs, an iterative image reconstruction algorithm is implemented. Computer simulation and experimental studies are conducted to investigate the computational efficiency and numerical accuracy of the developed algorithms.

Results: The GPU implementations improve the computational efficiency by factors of 1000, 125, and 250 for the FBP algorithm and the two pairs of projection/backprojection operators, respectively. Accurate images are reconstructed by use of the FBP and iterative image reconstruction algorithms from both computer-simulated and experimental data.

Conclusions: Parallelization strategies for 3D OAT image reconstruction are proposed for the first time. These GPU-based implementations significantly reduce the computational time for 3D image reconstruction, complementing our earlier work on 3D OAT iterative image reconstruction. © 2013 American Association of Physicists in Medicine. [<http://dx.doi.org/10.1118/1.4774361>]

Key words: Optoacoustic tomography, photoacoustic tomography, thermoacoustic tomography, graphics processing unit (GPU), compute unified device architecture (CUDA)

I. INTRODUCTION

Optoacoustic tomography (OAT), also known as photoacoustic computed tomography, is an emerging imaging modality that has great potential for a wide range of biomedical imaging applications.^{1–4} In OAT, a short laser pulse is employed to irradiate biological tissues. When the biological tissues absorb the optical energy, acoustic wave fields can be generated via the thermoacoustic effect. The acoustic wave fields propagate outward in three-dimensional (3D) space and are measured by use of ultrasonic transducers that are distributed outside the object. The goal of OAT is to obtain an estimate of the absorbed energy density map within the object from the measured acoustic signals. To accomplish this, an image reconstruction algorithm is required.

A variety of analytic image reconstruction algorithms have been proposed.^{5–8} These algorithms generally assume an idealized transducer model and an acoustically homogeneous medium. Also, since they are based on discretization of con-

tinuous reconstruction formulae, these algorithms require the acoustic pressure to be densely sampled over a surface that encloses the object to obtain an accurate reconstruction. To overcome these limitations, iterative image reconstruction algorithms have been proposed.^{9–24} Although the optoacoustic wave intrinsically propagates in 3D space, when applying to experimental data, most studies have employed two-dimensional (2D) imaging models by making certain assumptions on the transducer responses and/or the object structures.^{10,13,15,16,19,25} An important reason is because the computation required for 3D OAT image reconstruction is excessively burdensome. Therefore, acceleration of 3D image reconstruction will facilitate algorithm development and many applications including real-time 3D PACT.^{26,27}

A graphics processing unit (GPU) card is a specialized device specifically designed for parallel computations.²⁸ Compute unified device architecture (CUDA) is an extension of the C/FORTRAN language that provides a convenient programming platform to exploit the parallel computational power of

GPUs.²⁹ The CUDA-based parallel programming technique has been successfully applied to accelerate image reconstruction in mature imaging modalities such as x-ray computed tomography (CT)^{30–32} and magnetic resonance imaging (MRI).³³ In OAT, however, only a few works on utilization of GPUs to accelerate image reconstruction have been reported.^{20,34} For example, the k-wave toolbox employs the NVIDIA CUDA Fast Fourier Transform library (cuFFT) to accelerate the computation of 3D FFT.³⁴ Also a GPU-based sparse matrix-vector multiplication strategy has been applied to 3D OAT image reconstruction for the case that the system matrix is sparse and can be stored in memory.²⁰ However, there remains an important need to develop efficient implementations of OAT reconstruction algorithms for general applications in which the system matrix is too large to be stored.

In this work, we propose parallelization strategies, for use with GPUs, to accelerate 3D image reconstruction in OAT. Both filtered backprojection (FBP) and iterative image reconstruction algorithms are investigated. For use with iterative image reconstruction algorithms, we focus on the parallelization of projection and backprojection operators. Specifically, we develop two pairs of projection/backprojection operators that correspond to two distinct discrete-to-discrete (D-D) imaging models employed in OAT, namely, the interpolation-based and the spherical-voxel-based D-D imaging models. Note that our implementations of the backprojection operators compute the exact adjoint of the forward operators, and therefore the projector pairs are “matched.”³⁵

The remainder of the paper is organized as follows. In Sec. II, we briefly review OAT imaging models in their continuous and discrete forms. We propose GPU-based parallelization strategies in Sec. III. Numerical studies and results are described in Secs. IV and V, respectively. Finally, a brief discussion and summary of the proposed algorithms are provided in Sec. VI.

II. BACKGROUND

II.A. Continuous-to-continuous (C-C) imaging models and analytic image reconstruction algorithms

A C-C OAT imaging model neglects sampling effects and provides a mapping from the absorbed energy density function $A(\mathbf{r})$ to the induced acoustic pressure function $p(\mathbf{r}^s, t)$. Here, t is the temporal coordinate, $\mathbf{r} \in V$ and $\mathbf{r}^s \in S$ denote the locations within the object support V and on the measurement surface S , respectively. A canonical OAT C-C imaging model can be expressed as^{1,14,36}

$$p(\mathbf{r}^s, t) = \frac{\beta}{4\pi C_p} \int_V d\mathbf{r} A(\mathbf{r}) \frac{d}{dt} \frac{\delta\left(t - \frac{|\mathbf{r}^s - \mathbf{r}|}{c_0}\right)}{|\mathbf{r}^s - \mathbf{r}|} \equiv \mathcal{H}_{CC} A, \quad (1)$$

where $\delta(t)$ is the Dirac delta function, β , c_0 , and C_p denote the thermal coefficient of volume expansion, (constant) speed-of-sound, and the specific heat capacity of the medium at constant pressure, respectively. We introduce an operator notation $\mathcal{H}_{CC}$ to denote this C-C mapping.

Alternatively, Eq. (1) can be reformulated as the well-known spherical Radon transform (SRT)^{12,37}

$$g(\mathbf{r}^s, t) = \int_V d\mathbf{r} A(\mathbf{r}) \delta(c_0 t - |\mathbf{r}^s - \mathbf{r}|), \quad (2)$$

where the function $g(\mathbf{r}^s, t)$ is related to $p(\mathbf{r}^s, t)$ as

$$p(\mathbf{r}^s, t) = \frac{\beta}{4\pi C_p} \frac{\partial}{\partial t} \left(\frac{g(\mathbf{r}^s, t)}{t} \right). \quad (3)$$

The SRT model provides an intuitive interpretation of each value of $g(\mathbf{r}^s, t)$ as a surface integral of $A(\mathbf{r})$ over a sphere centered at $\mathbf{r}^s$ with radius tc_0 .

Based on C-C imaging models, a variety of analytic image reconstruction algorithms have been developed.^{5–8} For the case of a spherical measurement geometry, a FBP algorithm in its continuous form is given by⁶

$$A(\mathbf{r}) = -\frac{C_p}{2\pi\beta c_0^2 R^s} \int_S d\mathbf{r}^s \left[\frac{2p(\mathbf{r}^s, t)}{|\mathbf{r} - \mathbf{r}^s|} + \frac{1}{c_0} \frac{\partial p(\mathbf{r}^s, t)}{\partial t} \right]_{t=\frac{|\mathbf{r}-\mathbf{r}^s|}{c_0}}, \quad (4)$$

where R^s denotes the radius of the measurement surface S .

II.B. Discrete-to-discrete imaging models and iterative image reconstruction algorithms

When sampling effects are considered, an OAT system is properly described as a continuous-to-discrete (C-D) imaging model^{14,22,35,38,39}

$$[\mathbf{u}]_{qK+k} = h^e(t) *_t \frac{1}{S_q} \int_{S_q} d\mathbf{r}^s p(\mathbf{r}^s, t) \Big|_{t=k\Delta_t},$$

$$q = 0, 1, \dots, Q-1$$

$$k = 0, 1, \dots, K-1, \quad (5)$$

where Q and K denote the total numbers of transducers (indexed by q) and the time samples (indexed by k) respectively. S_q is the surface area of the q th transducer, which is assumed to be a subset of S ; $h^e(t)$ denotes the acousto-electric impulse response (EIR) of each transducer that, without loss of generality, is assumed to be identical for all transducers; “ $*_t$ ” denotes a linear convolution with respect to time coordinate; and Δ_t is the temporal sampling interval. The vector $\mathbf{u}$ represents the lexicographically ordered measured voltage signals whose $(qK + k)$ th element is denoted by $[\mathbf{u}]_{qK+k}$.

In order to apply iterative image reconstruction algorithms, a D-D imaging model is required, which necessitates the discretization of $A(\mathbf{r})$. The following N -dimensional representation of the object function can be employed^{35,38}

$$A(\mathbf{r}) \approx \sum_{n=0}^{N-1} [\boldsymbol{\alpha}]_n \psi_n(\mathbf{r}), \quad (6)$$

where $\boldsymbol{\alpha}$ is a coefficient vector whose n th element is denoted by $[\boldsymbol{\alpha}]_n$ and $\psi_n(\mathbf{r})$ is the expansion function. On substitution from Eq. (6) into Eq. (5), where $p(\mathbf{r}^s, t)$ is defined by Eq. (1),

one obtains a D-D mapping from α to $\mathbf{u}$, expressed as

$$\mathbf{u} \approx \mathbf{H}\alpha, \quad (7)$$

where each element of the matrix $\mathbf{H}$ is defined as

$$[\mathbf{H}]_{qK+k,n} = \left[h^e *_t \frac{1}{S_q} \int_{S_q} d\mathbf{r}^s \mathcal{H}_{CC} \psi_n \right]_{t=k\Delta_t}. \quad (8)$$

Here, $\mathbf{H}$ is the D-D imaging operator also known as system matrix or projection operator. Note that the “ $\approx$ ” in Eq. (7) is due to the use of the finite-dimensional representation of the object function [i.e., Eq. (6)]. No additional approximations have been introduced.

Below we describe two types of D-D imaging models that have been employed in OAT:^{14,22,39} the interpolation-based imaging model and the spherical-voxel-based imaging model. The quantities $\mathbf{u}$, $\mathbf{H}$, and α (or ψ_n) in the two models will

be distinguished by the subscripts (or superscripts) “int” and “sph,” respectively.

II.B.1. Interpolation-based D-D imaging model

The interpolation-based D-D imaging model defines the coefficient vector as samples of the object function on the nodes of a uniform Cartesian grid

$$[\alpha_{\text{int}}]_n = \int_V d\mathbf{r} \delta(\mathbf{r} - \mathbf{r}_n) A(\mathbf{r}), \quad n = 0, 1, \dots, N-1, \quad (9)$$

where $\mathbf{r}_n = (x_n, y_n, z_n)^T$ specifies the location of the n th node of the uniform Cartesian grid. The definition of the expansion function depends on the choice of interpolation method.¹⁹ If a trilinear interpolation method is employed, the expansion function can be expressed as⁴⁰

$$\psi_n^{\text{int}}(\mathbf{r}) = \begin{cases} \left(1 - \frac{|x-x_n|}{\Delta_s}\right) \left(1 - \frac{|y-y_n|}{\Delta_s}\right) \left(1 - \frac{|z-z_n|}{\Delta_s}\right), & \text{if } |x-x_n|, |y-y_n|, |z-z_n| \leq \Delta_s, \\ 0, & \text{otherwise} \end{cases}, \quad (10)$$

where Δ_s is the distance between two neighboring grid points.

In principle, the interpolation-based D-D imaging model can be constructed by substitution from Eqs. (9) and (10) to Eq. (8). In practice, however, implementation of the surface integral over S_q is difficult for the choice of expansion functions in Eq. (10). Also, implementations of the temporal convolution and $\mathcal{H}_{CC} \psi_n^{\text{int}}$ usually require extra discretization procedures. Therefore, utilization of the interpolation-based D-D model commonly assumes the transducers to be point-like. In this case, the implementation of $\mathbf{H}_{\text{int}}$ is decomposed as a three-step operation

$$\mathbf{u}_{\text{int}} = \mathbf{H}_{\text{int}} \alpha_{\text{int}} \equiv \mathbf{H}^e \mathbf{D} \mathbf{G} \alpha_{\text{int}}, \quad (11)$$

where $\mathbf{G}$, $\mathbf{D}$, and $\mathbf{H}^e$ are discrete approximations of the SRT [Eq. (2)], the differential operator [Eq. (3)], and the operator that implements a temporal convolution with EIR, respectively. We implemented $\mathbf{G}$ in a way^{12,41,42} that is similar to the “ray-driven” implementation of Radon transform in x-ray CT,⁴⁰ i.e., for each data sample, we accumulated the contributions from the voxels that resided on the spherical shell specified by the data sample. By use of Eqs. (2), (6), (9) and (10), one obtains

$$[\mathbf{G} \alpha_{\text{int}}]_{qK+k} = \Delta_s^2 \sum_{n=0}^{N-1} [\alpha_{\text{int}}]_n \sum_{i=0}^{N_i-1} \sum_{j=0}^{N_j-1} \psi_n^{\text{int}}(\mathbf{r}_{k,i,j}) \equiv [\mathbf{g}]_{qK+k}, \quad (12)$$

where $[\mathbf{g}]_{qK+k} \approx g(\mathbf{r}_q^s, t)|_{t=k\Delta_t}$ with $\mathbf{r}_q^s$ specifying the location of the q th point-like transducer, and N_i and N_j denote the numbers of divisions over the two angular coordinates of a local spherical coordinate system shown in Fig. 1(b). A derivation of Eq. (12) is provided in the Appendix. The differential

FIG. 1. (a) Schematic of the 3D OAT scanning geometry. (b) Schematic of the local coordinate system for the implementation of interpolation-based D-D imaging model.

operator in Eq. (3) is approximated as

$$[\mathbf{D}\mathbf{g}]_{qK+k} = \frac{\beta}{8\pi C_p \Delta_t^2} \left(\frac{[\mathbf{g}]_{qK+k+1}}{k+1} - \frac{[\mathbf{g}]_{qK+k-1}}{k-1} \right) \equiv [\mathbf{p}_{\text{int}}]_{qK+k}, \quad (13)$$

where $[\mathbf{p}_{\text{int}}]_{qK+k} \approx p(\mathbf{r}_q^s, t)|_{t=k\Delta_t}$. Finally, the continuous temporal convolution is approximated by a discrete linear convolution as⁴³

$$[\mathbf{H}^e \mathbf{p}_{\text{int}}]_{qK+k} = \sum_{\kappa=0}^{K-1} [\mathbf{h}^e]_{k-1-\kappa} [\mathbf{p}_{\text{int}}]_{qK+\kappa} \equiv [\mathbf{u}_{\text{int}}]_{qK+k}, \quad (14)$$

where $[\mathbf{h}^e]_k = \Delta_t h^e(t)|_{t=k\Delta_t}$.

II.B.2. Spherical-voxel-based D-D imaging model

The spherical-voxel-based imaging model is also widely employed in OAT.^{9,11,14,22,44} It employs the expansion functions

$$\psi_n^{\text{sph}}(\mathbf{r}) = \begin{cases} 1, & \text{if } |\mathbf{r} - \mathbf{r}_n| \leq \Delta_s/2 \\ 0, & \text{otherwise} \end{cases}, \quad (15)$$

where $\mathbf{r}_n$ is defined as in Eq. (9). The n th expansion function $\psi_n^{\text{sph}}(\mathbf{r})$ is a uniform sphere that is inscribed by the n th cuboid of a Cartesian grid. The n th component of the coefficient vector α_{sph} is defined as

$$[\alpha_{\text{sph}}]_n = \frac{V_{\text{cube}}}{V_{\text{sph}}} \int_V d\mathbf{r} \psi_n^{\text{sph}}(\mathbf{r}) A(\mathbf{r}), \quad (16)$$

where V_{cube} and V_{sph} are the volumes of a cubic voxel of dimension Δ_s and of a spherical voxel of radius $\Delta_s/2$, respectively.

Unlike the interpolation-based imaging model, by use of the expansion functions defined in Eq. (15), the surface integral over S_q and $\mathcal{H}_{\text{CC}} \psi_n^{\text{sph}}$ in Eq. (8) can be converted to a temporal convolution and calculated analytically.^{14,22} To avoid utilizing excessively high sampling rates to mitigate aliasing, the spherical-voxel-based imaging model can be conveniently implemented in the temporal frequency domain as²²

$$[\tilde{\mathbf{u}}_{\text{sph}}]_{qL+l} = \tilde{p}_0(f) \sum_{n=0}^{N-1} [\alpha_{\text{sph}}]_n \frac{1}{S_q} \tilde{h}_q^s(\mathbf{r}_n, f) \Big|_{f=l\Delta_f}, \quad (17)$$

for $l = 0, 1, \dots, L-1$,

where Δ_f is the frequency sampling interval, and L denotes the total number of temporal-frequency samples indexed by l . A derivation of Eq. (17) can be found in Ref. 22. The function $\tilde{h}_q^s(\mathbf{r}_n, f)$ represents the temporal Fourier transform of the spatial impulse response (SIR) of the q th transducer for the source located at $\mathbf{r}_n$, expressed as

$$\tilde{h}_q(\mathbf{r}_n, f) = \int_{S_q} d\mathbf{r}^s \frac{\exp(-j2\pi f \frac{|\mathbf{r}^s - \mathbf{r}_n|}{c_0})}{2\pi |\mathbf{r}^s - \mathbf{r}_n|}. \quad (18)$$

Also, $\tilde{p}_0(f)$ is defined as

$$\tilde{p}_0(f) = -j \frac{\beta c_0^3}{C_p f} \left[\frac{\Delta_s}{2c_0} \cos\left(\frac{\pi f \Delta_s}{c_0}\right) - \frac{1}{2\pi f} \sin\left(\frac{\pi f \Delta_s}{c_0}\right) \right] \tilde{h}^e(f), \quad (19)$$

where $\tilde{h}^e(f)$ is the EIR in temporal-frequency domain. In summary, the imaging model can be expressed in matrix form as

$$\tilde{\mathbf{u}}_{\text{sph}} = \mathbf{H}_{\text{sph}} \alpha_{\text{sph}}. \quad (20)$$

II.B.3. Adjoints of the system matrices

Iterative image reconstruction algorithms employ numerical implementations of the projection operator, i.e., the system matrix $\mathbf{H}$, as well as its adjoint, denoted by $\mathbf{H}^\dagger$.⁴⁵ The adjoint is also referred to as the backprojection operator. Note that for most practical applications, $\mathbf{H}$ and $\mathbf{H}^\dagger$ are too large to be stored in the random access memory of currently available computers. Therefore, in practice, the actions of $\mathbf{H}$ and $\mathbf{H}^\dagger$ are almost always calculated on the fly. The same strategy was adopted in this work.

According to the definition of the adjoint operator,^{35,43} $\mathbf{H}_{\text{int}}^\dagger = \mathbf{G}^\dagger \mathbf{D}^\dagger \mathbf{H}^e$, where

$$[\mathbf{H}^e \mathbf{u}_{\text{int}}]_{qK+k} = \sum_{\kappa=0}^{K-1} [\mathbf{h}^e]_{\kappa-1-k} [\mathbf{u}_{\text{int}}]_{qK+\kappa} \equiv [\mathbf{p}'_{\text{int}}]_{qK+k}, \quad (21)$$

$$[\mathbf{D}^\dagger \mathbf{p}'_{\text{int}}]_{qK+k} = \frac{\beta}{8\pi C_p \Delta_t^2 k} ([\mathbf{p}'_{\text{int}}]_{qK+k-1} - [\mathbf{p}'_{\text{int}}]_{qK+k+1}) \equiv [\mathbf{g}']_{qK+k}, \quad (22)$$

and

$$[\mathbf{G}^\dagger \mathbf{g}']_n = \Delta_s^2 \sum_{q=0}^{Q-1} \sum_{k=0}^{K-1} [\mathbf{g}']_{qK+k} \sum_{i=0}^{N_i-1} \sum_{j=0}^{N_j-1} \psi_n^{\text{int}}(\mathbf{r}_{k,i,j}) \equiv [\alpha'_{\text{int}}]_n. \quad (23)$$

It can also be verified that the adjoint operator $\mathbf{H}_{\text{sph}}^\dagger$ is given by

$$[\mathbf{H}_{\text{sph}}^\dagger \tilde{\mathbf{u}}_{\text{sph}}]_n = \sum_{q=0}^{Q-1} \sum_{l=0}^{L-1} [\tilde{\mathbf{u}}_{\text{sph}}]_{qL+l} \tilde{p}_0^*(f) \tilde{h}_q^*(\mathbf{r}_n, f) \Big|_{f=l\Delta_f}, \quad (24)$$

where the superscript “*” denotes the complex conjugate. Unlike the unmatched backprojection operators⁴⁶ that are obtained by discretization of the continuous adjoint operator, utilization of the exact adjoint operator facilitates the convergence of iterative image reconstruction algorithms.

II.C. GPU architecture and CUDA programming

The key features of GPU architecture and the basics of CUDA programming are briefly summarized in this section. We refer the readers to Refs. 28 and 29 for additional details.

A GPU card contains multiple streaming multiprocessors. Each streaming multiprocessor is configured with multiple processor cores. For example, the Tesla C1060 possesses 30 streaming multiprocessors with 8 processor cores on each; and the Tesla C2050 possesses 14 streaming multiprocessors with 32 processor cores on each.²⁸ The processor cores in each multiprocessor execute the same instruction on different pieces of data, which is referred to as “single instruction,

multiple data" (SIMD) model of parallel programming. In order to fully exploit the computing power of GPUs, one of the major challenges is to design a parallelization strategy fitting in the SIMD framework such that the largest number of processor cores can execute the computation simultaneously.²⁹

A GPU card has six types of memory that have varying capacities and different access rules and efficiencies: (1) Registers are assigned for each thread and have the fastest access. (2) Shared memory is assigned for each block and can be efficiently accessed by all threads in the block if designed appropriately. (3) Constant memory is read-only and can be accessed by all threads efficiently. (4) Texture memory is also read-only and is optimized for interpolation operations. (5) Global memory has the slowest access that takes hundred times more clock cycles than does the computation of basic arithmetic operations. (6) Local memory is assigned for each thread but has a slow access as does the global memory. Therefore, an efficient GPU-based implementation in general requires a limited number of global and local memory access.

CUDA is a platform and programming model developed by NVIDIA that includes a collection of functions and keywords to exploit the parallel computing power of GPUs.²⁹ A CUDA parallel program is composed of a host program and kernels. The host program is executed by central processing units (CPUs) and launches the kernels, which are custom-designed functions executed by GPUs. A general parallel programming strategy is to launch multiple instances of a kernel and to run the multiple instances concurrently on GPUs. In CUDA, each instance of the kernel is named as a thread and processes only a portion of the data. A hierarchy of threads is employed: Threads are grouped into blocks, and blocks are grouped into a grid. Therefore, each thread is specified by a multi-index containing a block index and a thread index within the block.

III. GPU-ACCELERATED RECONSTRUCTION ALGORITHMS

In this section, we propose GPU-based parallelization strategies for the FBP algorithm and the projection/backprojection operations corresponding to the interpolation-based and the spherical-voxel-based D-D imaging models.

III.A. Measurement geometry

We employed a spherical measurement geometry shown in Fig. 1(a). The measurement sphere was of radius R^s centered at the origin of the Cartesian coordinate system (or the equivalent spherical coordinate system). The polar angle $\theta^s \in [0, \pi]$ was equally divided with interval $\Delta\theta^s = \pi/N_r$, starting from $\theta_{\min}^s$. At each polar angle, a ring on the sphere that was parallel to the plane $z = 0$ can be specified, resulting N_r rings. On each ring, N_v ultrasonic transducers were assumed to be uniformly distributed with azimuth angle interval $\Delta\phi^s = 2\pi/N_v$. Hereafter, each azimuth angle will be referred to as a tomographic view. At each view, we assumed that N_t temporal samples were acquired and the first sample corresponded to time instance $t_{\min}$. For implementations in temporal-frequency do-

main, we assumed that N_f temporal-frequency samples were available and the first sample corresponded to $f_{\min}$. The region to be reconstructed was a rectangular cuboid whose edges were parallel to the axes of the coordinate system and the left-bottom-back vertex was located at $(x_{\min}, y_{\min}, z_{\min})$. The numbers of voxels along the three coordinates will be denoted by N_x , N_y , and N_z , respectively, totally $N = N_x N_y N_z$ voxels. We also assumed the cuboid was contained in another sphere of radius R that was concentric with the measurement sphere shown in Fig. 1(b).

III.B. Implementation of the FBP algorithm

Central processing unit-based implementations of continuous FBP formulae have been described in Refs. 5–8. Though the discretization methods vary, in general, three approximations have to be employed. First, the first-order derivative term $\partial p(\mathbf{r}^s, t)/\partial t$ has to be approximated by a difference scheme up to certain order.⁴⁷ Second, the measurement sphere has to be divided into small patches, and the surface integral has to be approximated by a summation of the area of every patch weighted by the effective value of the integrand on the patch. Finally, the value of the integrand at an arbitrary time instance $t = |\mathbf{r}^s - \mathbf{r}|/c_0$ has to be approximated by certain interpolation method.

In this study, we approximated the surface integral by use of the trapezoidal rule. As described earlier, the spherical surface was divided into $N_r N_v$ patches. For the transducer indexed by q that was located at $\mathbf{r}_q^s = (R^s, \theta_q^s, \phi_q^s)$, the area of the patch was approximated by $(R^s)^2 \Delta\theta^s \Delta\phi^s \sin\theta_q^s$. The value at time instance $t = |\mathbf{r}_q^s - \mathbf{r}_n|/c_0$ was approximated by the linear interpolation from its two neighboring samples as

$$p(\mathbf{r}_q^s, t) \Big|_{t=\frac{|\mathbf{r}_q^s - \mathbf{r}_n|}{c_0}} \approx (k+1-\tilde{k})[\mathbf{p}]_{qK+k} + (\tilde{k}-k)[\mathbf{p}]_{qK+k+1}, \quad (25)$$

where $\tilde{k} = (|\mathbf{r}_q^s - \mathbf{r}_n|/c_0 - t_{\min})/\Delta t$, and k is the integer part of $\tilde{k}$. Here, $\mathbf{p}$ is a vector of lexicographically ordered samples of the pressure function $p(\mathbf{r}^s, t)$, which is estimated from the measured voltage data vector $\mathbf{u}$. Also, the first-order derivative term was approximated by

$$\frac{\partial}{\partial t} p(\mathbf{r}_q^s, t) \Big|_{t=\frac{|\mathbf{r}_q^s - \mathbf{r}_n|}{c_0}} \approx \frac{1}{\Delta t} ([\mathbf{p}]_{qK+k+1} - [\mathbf{p}]_{qK+k}). \quad (26)$$

By use of these three numerical approximations, the discretized FBP formula was expressed as

$$\begin{aligned} [\hat{\alpha}_{\text{fbp}}]_n = & -\frac{C_p R^s \Delta\theta^s \Delta\phi^s}{\pi \beta c_0^3 \Delta t} \sum_{n_r=0}^{N_r-1} \sin\theta_q^s \\ & \times \sum_{n_v=0}^{N_v-1} \left\{ \left(1.5 - \frac{k + t_{\min}/\Delta t}{\tilde{k} + t_{\min}/\Delta t} \right) [\mathbf{p}]_{qK+k+1} \right. \\ & \left. + \left(\frac{k+1 + t_{\min}/\Delta t}{\tilde{k} + t_{\min}/\Delta t} - 1.5 \right) [\mathbf{p}]_{qK+k} \right\}. \end{aligned} \quad (27)$$

Unlike the implementations of FBP formulas in x-ray cone beam CT,^{31,32} we combined the filter and the linear

ALGORITHM I. Implementation of the FBP algorithm (on host).

Input: $\mathbf{p}$ **Output:** $\hat{\alpha}_{\text{fbp}}$

- 1: $w = -C_p R^s \Delta_{\theta^s} \Delta_{\phi^s} / (\pi \beta c_0^3 \Delta_t)$ {Precalculate the common coefficient}
- 2: $\mathbf{T}_{\mathbf{p}} \leftarrow \mathbf{p}$ {Bound data to texture memory}
- 3: $\mathbf{K}_{\text{fbp}} \langle \langle \langle N_y, N_x \rangle, (N_z, 1, 1) \rangle \rangle (\omega, \mathbf{D}_{\hat{\alpha}_{\text{fbp}}})$
- 4: $\hat{\alpha}_{\text{fbp}} \leftarrow \mathbf{D}_{\hat{\alpha}_{\text{fbp}}}$ {Copy data from global memory to host}

interpolation. This reduced the number of visits to the global memory in the GPU implementation described below.

We implemented the FBP formula in a way that is similar to the “pixel-driven” implementation in x-ray CT,³² i.e., we assigned each thread to execute the two accumulative summations in Eq. (27) for each voxel. We bound the pressure data $\mathbf{p}$ to texture memory because it is cached and has a faster accessing rate. Therefore, our implementation only requires access to texture memory twice and to global memory once. The pseudocodes are provided in Algorithms 1 and 2 for the host part and the device part, respectively. Note that the pseudocodes do not intend to be always optimal because the performance of the codes could depend on the dimensions of $\mathbf{p}$ and $\hat{\alpha}_{\text{fbp}}$. For example, we set the block size to be $(N_z, 1, 1)$ because for our applications, N_z was bigger than N_x and N_y and smaller than the limit number of threads that a block can support (i.e., 1024 for the NVIDIA Tesla C2050). If the values of N_x , N_y , and N_z change, we may need to redesign the dimensions of the grid and blocks. However, the general SIMD parallelization strategy remains.

III.C. Implementation of $\mathbf{H}_{\text{int}}$ and $\mathbf{H}_{\text{int}}^\dagger$

The forward projection operation $\mathbf{H}_{\text{int}}\alpha_{\text{int}}$ is composed of three consecutive operations $\mathbf{g} = \mathbf{G}\alpha_{\text{int}}$, $\mathbf{p}_{\text{int}} = \mathbf{D}\mathbf{g}$, and $\mathbf{u}_{\text{int}} = \mathbf{H}^e \mathbf{p}_{\text{int}}$ that are defined in Eqs. (12)–(14), respectively. Both the difference operator $\mathbf{D}$ and the one-dimensional (1D) convolution $\mathbf{H}^e$ have low computational complexities while the SRT operator $\mathbf{G}$ is computationally burdensome. Hence, we

ALGORITHM III. Implementation of $\mathbf{g} = \mathbf{G}\alpha_{\text{int}}$ (on host)**Input:** α_{int} **Output:** $\mathbf{g}$

- 1: $\mathbf{T}_{\alpha_{\text{int}}} \leftarrow \alpha_{\text{int}}$ {Bound data to texture memory}
- 2: $\mathbf{K}_{\text{srt}} \langle \langle \langle N_v, N_t \rangle, (N_r, 1, 1) \rangle \rangle (\mathbf{D}_{\mathbf{g}})$
- 3: $\mathbf{g} \leftarrow \mathbf{D}_{\mathbf{g}}$ {Copy data from global memory to host}

developed the GPU-based implementation of $\mathbf{G}$ while leaving $\mathbf{D}$ and $\mathbf{H}^e$ to be implemented by CPUs.

The SRT in OAT shares many features with the Radon transform in x-ray CT. Thus, our GPU-based implementation is closely related to the implementations of Radon transform that have been optimized for x-ray CT.^{30–32} The surface integral was approximated according to the trapezoidal rule. First, the integral surface was divided into small patches, which is described in the Appendix. Second, each patch was assigned an effective value of the object function by trilinear interpolation. The trilinear interpolation was calculated by use of the texture memory of GPUs that is specifically designed for interpolation. Finally, GPU threads accumulated the areas of patches weighted by the effective values of the object function and wrote the final results to global memory. The pseudocodes for implementation of $\mathbf{G}$ are provided in Algorithms 3 and 4 for the host part and the device part, respectively. Note that we employed the “one-level”-strategy,³² i.e., each thread calculates one data sample. Higher level strategies have been proposed to improve the performance by assigning each block to calculate multiple data samples,³² which, however, caused many thread idles in OAT mainly because the amount of computation required to calculate a data sample varies largely among samples for SRT.

Implementation of the backprojection operator $\mathbf{H}_{\text{int}}^\dagger$ was very similar to the implementation of $\mathbf{H}_{\text{int}}$. The operators $\mathbf{D}^\dagger$ and $\mathbf{H}^{e\dagger}$ were calculated on CPUs while $\mathbf{G}^\dagger$ was calculated by use of GPUs. The pseudocodes are provided in Algorithms 5 and 6. We made use of the CUDA function “atomicAdd” to add weights to global memory from each thread.

ALGORITHM II. Implementation of kernel $\mathbf{K}_{\text{fbp}} \langle \langle \langle N_y, N_x \rangle, (N_z, 1, 1) \rangle \rangle$.**Input:** $\omega, \mathbf{T}_{\mathbf{p}}, \mathbf{D}_{\hat{\alpha}_{\text{fbp}}}$ **Output:** $\mathbf{D}_{\hat{\alpha}_{\text{fbp}}}$

- 1: $x = (\text{blockIdx.y})\Delta_s + x_{\min}; y = (\text{blockIdx.x})\Delta_s + y_{\min}; z = (\text{threadIdx.x})\Delta_s + z_{\min}$
- 2: $\Sigma = 0$
- 3: **for** $n_r = 0$ **to** $N_r - 1$ **do**
- 4: $\theta^s = n_r \Delta_{\theta^s} + \theta_{\min}^s; z^s = R^s \cos \theta^s; r^s = R^s \sin \theta^s; w' = w \sin \theta^s$
- 5: **for** $n_v = 0$ **to** $N_v - 1$ **do**
- 6: $\phi^s = n_v \Delta_{\phi^s} + \phi_{\min}^s; x^s = r^s \cos \phi^s; y^s = r^s \sin \phi^s$
- 7: $\tilde{r} = ((x - x^s)^2 + (y - y^s)^2 + (z - z^s)^2)^{1/2}$
- 8: $t_n = (\tilde{r}/c_0 - t_{\min})/\Delta_t; n_t = \text{floor}(t_n)$
- 9: $\Sigma += w' \{[(n_t \Delta_t + t_{\min})/(t_n \Delta_t + t_{\min}) - 1.5] \mathbf{T}_{\mathbf{p}}[n_r][n_v][n_t] + [1.5 - (n_t \Delta_t + t_{\min})/(t_n \Delta_t + t_{\min})] \mathbf{T}_{\mathbf{p}}[n_r][n_v][n_t + 1]\}$ {Fetch data from texture memory}
- 10: **end for**
- 11: **end for**
- 12: $\mathbf{D}_{\hat{\alpha}_{\text{fbp}}}[\text{blockIdx.y}][\text{blockIdx.x}][\text{threadIdx.x}] = \Sigma$

ALGORITHM IV. Implementation of kernel $K_{\text{srt}} \langle \langle (N_v, N_t), (N_r, 1, 1) \rangle \rangle$.

Input: $D_g, T_{\alpha_{\text{int}}}$

Output: D_g

```

1:  $\bar{r} = (\text{blockIdx.y})c_0\Delta_t + c_0t_{\text{min}}; \quad \theta^s = (\text{threadIdx.x})\Delta\theta^s + \theta_{\text{min}}^s; \quad \phi^s = (\text{blockIdx.x})\Delta\phi^s + \phi_{\text{min}}^s$ 
2:  $\theta'_{\text{max}} = \arccos(((R^s)^2 + \bar{r}^2 - R^2)/(2\bar{r}R^s))$ 
3:  $\Sigma = 0; \quad \theta' = \theta'_{\text{max}}$ 
4: while  $\theta' > 0$  do
5:    $z' = \bar{r} \cos \theta'; \quad r' = \bar{r} \sin \theta'; \quad \phi' = 0$ 
6:   while  $\phi' < 2\pi$  do
7:      $x' = r' \cos \phi'; \quad y' = r' \sin \phi'$ 
8:      $x = -x' \sin \theta' - (z' - R^s) \cos \theta'; \quad y = y'; \quad z = x' \cos \theta' - (z' - R^s) \sin \theta' \quad \{\text{Convert to global coordinate system}\}$ 
9:      $x_n = (x - x_{\text{min}})/\Delta_s; \quad y_n = (y - y_{\text{min}})/\Delta_s; \quad z_n = (z - z_{\text{min}})/\Delta_s$ 
10:     $\Sigma += \text{tex3D}(x_n, y_n, z_n) \quad \{\text{Tri-linear interpolation}\}$ 
11:     $\phi' = \phi' + \Delta_s/r'$ 
12:   end while
13:    $\theta' = \theta' - \Delta_s/\bar{r}$ 
14: end while
15:  $D_g[\text{threadIdx.x}][\text{blockIdx.x}][\text{blockIdx.y}] = \Sigma \Delta_s^2$ 

```

III.D. Implementation of H_{sph} and $H_{\text{sph}}^\dagger$

Implementation of the forward projection operation for the spherical-voxel-based imaging model is distinct from that of the interpolation-based model. The major difference is that calculation of each element of the data vector for the spherical-voxel-based imaging model requires the accumulation of the contributions from all voxels because the model is expressed in the temporal frequency domain. Because of this, the amount of computation required to calculate each data sample in the spherical-voxel-based imaging model is almost identical, simplifying the parallelization strategy.

We proposed a parallelization strategy that was inspired by one applied in advanced MRI reconstruction³³ and is summarized as follows. Discrete samples of $\tilde{p}_0(f)$ defined in Eq. (19) were precalculated and stored as a vector $\tilde{\mathbf{p}}_0$ in constant memory. Because the size of the input vector α_{sph} is often too large to fit in the constant memory, we divided α_{sph} into subvectors that matched the capacity of the constant memory. We employed a CPU loop to copy every subvector sequentially to the constant memory and call the GPU kernel function to accumulate a partial summation. The major advantage of this design is that the total number of global memory visits to calculate one data sample is reduced to the number of subvectors.

Implementation of the projection operator for the spherical-voxel-based imaging model generally involves more arithmetic operations than does the interpolation-based imaging model. Moreover, the spherical-voxel-based imag-

ing model has been employed to compensate for the finite aperture size effect of transducers,^{14,22} which makes the computation even more burdensome. Because of this, we further developed an implementation that employed multiple GPUs. The pseudocodes of the projection operation are provided in Algorithms 7–9. We created N_{pth} pthreads on CPUs by use of the “pthread.h” library. Here, we denote the threads on CPUs by “pthread” to distinguish from threads on GPUs. We divided the input vector α_{sph} into N_{pth} subvectors (denoted by α_{pth} ’s) of equal size and declared an output vector $\tilde{\mathbf{u}}'_{\text{sph}}$ of dimension $N_{\text{pth}}N_rN_vN_f$. By calling the pthread function “fwd_pthread,” N_{sph} pthreads simultaneously calculated the projection. Each pthread projected an α_{pth} to a partial voltage data vector $\tilde{\mathbf{u}}'_{\text{pth}}$ that filled in the larger vector $\tilde{\mathbf{u}}'_{\text{sph}}$. Once all pthreads finished filling their $\tilde{\mathbf{u}}'_{\text{pth}}$ into $\tilde{\mathbf{u}}'_{\text{sph}}$, the projection data $\tilde{\mathbf{u}}_{\text{sph}}$ were obtained by a summation of the N_{pth} $\tilde{\mathbf{u}}'_{\text{pth}}$ ’s.

Implementation of the backprojection operator was similar except the dividing and looping were over the vector $\tilde{\mathbf{u}}_{\text{sph}}$ instead of α_{sph} . The pseudocodes for the backprojection operation are provided in Algorithms 10–12.

IV. DESCRIPTIONS OF COMPUTER SIMULATION AND EXPERIMENTAL STUDIES

The computational efficiency and accuracy of the proposed GPU-based implementations of the FBP algorithm and projection/backprojection operators for use with iterative image reconstruction algorithms were quantified in computer simulation and experimental OAT imaging studies.

IV.A. Computer-simulation studies

IV.A.1. Numerical phantom

The numerical phantom consisted of nine uniform spheres that were blurred by a 3D Gaussian kernel possessing a full width at half maximum (FWHM) of 0.77 mm. The

ALGORITHM V. Implementation of $\alpha'_{\text{int}} = \mathbf{G}^\dagger \mathbf{g}'$ (on host).

Input: $\mathbf{g}'$

Output: α'_{int}

```

1:  $T_{\mathbf{g}'} \leftarrow \mathbf{g}' \quad \{\text{Bound data to texture memory}\}$ 
2:  $K_{\text{srtT}} \langle \langle (N_v, N_t), (N_r, 1, 1) \rangle \rangle (D_{\alpha'_{\text{int}}})$ 
3:  $\alpha'_{\text{int}} \leftarrow D_{\alpha'_{\text{int}}} \quad \{\text{Copy data from global memory to host}\}$ 

```

ALGORITHM VI. Implementation of kernel $K_{\text{srtT}} \langle \langle (N_v, N_t), (N_r, 1, 1) \rangle \rangle$.

Input: $D_{\alpha'_{\text{int}}}, T_{\mathbf{g}'_{\text{int}}}$

Output: $D_{\alpha'_{\text{int}}}$

```

1:  $\bar{t} = (\text{blockIdx.y})c_0\Delta_t + c_0t_{\min}; \quad \theta^s = (\text{threadIdx.x})\Delta\theta^s + \theta^s_{\min}; \quad \phi^s = (\text{blockIdx.x})\Delta\phi^s + \phi^s_{\min}$ 
2:  $\theta'_{\max} = \arccos(((R^s)^2 + \bar{t}^2 - R^2)/(2\bar{t}R^s)); \quad \theta' = \theta'_{\max}$ 
3: while  $\theta' > 0$  do
4:    $z' = \bar{t} \cos \theta'; \quad r' = \bar{t} \sin \theta'; \quad \phi' = 0$ 
5:   while  $\phi' < 2\pi$  do
6:      $x' = r' \cos \phi'; \quad y' = r' \sin \phi'$ 
7:      $x = -x' \sin \theta - (z' - R') \cos \theta; \quad y = y'; \quad z = x' \cos \theta - (z' - R') \sin \theta$  {Convert to global coordinate system}
8:      $x_n = (x - x_{\min})/\Delta_s; \quad y_n = (y - y_{\min})/\Delta_s; \quad z_n = (z - z_{\min})/\Delta_s$ 
9:      $n_x = \text{floor}(x_n); \quad n_y = \text{floor}(y_n); \quad n_z = \text{floor}(z_n)$ 
10:     $D_{\alpha'_{\text{int}}}[n_z][n_y][n_x] += \Delta_s^2(n_x + 1 - x_n)(n_y + 1 - y_n)(n_z + 1 - z_n)T_{\mathbf{g}'_{\text{int}}}[\text{threadIdx.x}][\text{blockIdx.x}][\text{blockIdx.y}]$  {Add weights to one of the eight neighboring nodes by use of 'atomicAdd'; Repeat this operation for all other seven neighboring nodes}
11:     $\phi' = \phi' + \Delta_s/r'$ 
12:  end while
13:   $\theta' = \theta' - \Delta_s/\bar{t}$ 
14: end while

```

phantom was contained within a cuboid of size $29.4 \times 29.4 \times 61.6 \text{ mm}^3$. A 2D image corresponding to the plane $y = 0$ through the phantom is shown in Fig. 2(a).

IV.A.2. Simulated projection data

The measurement surface was a sphere of radius $R^s = 65 \text{ mm}$ corresponding to an existing OAT imaging system.^{22,48} As described in Sec. III, ideal point-like transducers were uniformly distributed over 128 rings and 90 tomographic views. The 128 rings covered the full π polar angle, i.e., $\theta^s_{\min}$

$= \pi/256$, while the 90 views covered the full 2π azimuth angle. The speed of sound was set at $c_0 = 1.54 \text{ mm}/\mu\text{s}$. We selected the Grüneisen coefficient as $\Gamma = \beta c_0^2/C_p = 2,000$ of arbitrary units (a.u.). For each transducer, we analytically calculated 1022 temporal samples of the pressure function at the sampling rate of $f_{\text{sam}} = 20 \text{ MHz}$ by use of Eq. (1). Because we employed a smooth object function, the pressure data were calculated by the following two steps: First, we calculated temporal samples of pressure function $p_{\text{us}}(\mathbf{r}^s, t)$ that corresponds to the nine uniform spheres by^{1,36}

$$p_{\text{us}}(\mathbf{r}^s, t)|_{t=k\Delta_t} = \sum_{i=0}^8 \begin{cases} A_i \left[-\frac{\beta c_0^3}{C_p |\mathbf{r}^s - \mathbf{r}_i|} t + \frac{\beta c_0^2}{2} \right]_{t=k\Delta_t}, & \text{if } |c_0 k \Delta_t - |\mathbf{r}^s - \mathbf{r}_i|| \leq R_i, \\ 0, & \text{otherwise} \end{cases}, \quad (28)$$

where $\mathbf{r}_i$, R_i , and A_i denote the center location, the radius, and the absorbed energy density of the i th sphere, respectively. Subsequently, we convolved $p_{\text{us}}(\mathbf{r}^s, t)$ with a 1D Gaussian kernel with FWHM = $0.5 \mu\text{s}$ (Ref. 49) to produce the pressure data. From the simulated pressure data, we calculated the temporal-frequency spectrum by use of FFT, from which we created an alternative data vector that contained 511 frequency components occupying (0, 5] MHz for each transducer. The simulated projection data in either the time domain or the temporal frequency domain will hereafter be referred to as “128 × 90”-data. By undersampling the “128 × 90”-data uniformly over rings and tomographic views, we created three subsets that contained varying number of transducers. These data sets will be referred to as “64 × 90”-data, “64 × 45”-data, and “32 × 45”-data, where the two numbers specify the number of rings and the number of tomographic views, respectively.

ALGORITHM VII. Implementation of $\tilde{\mathbf{u}}_{\text{sph}} = \mathbf{H}_{\text{sph}}\alpha_{\text{sph}}$ (on host).

Input: $\alpha_{\text{sph}}, \tilde{\mathbf{p}}_0$

Output: $\tilde{\mathbf{u}}_{\text{sph}}$

```

1: for  $n_{\text{pth}} = 0$  to  $N_{\text{pth}} - 1$  do
2:    $\text{parm\_fwdarg}[n_{\text{pth}}].n_{\text{pth}} = n_{\text{pth}}$ 
3:    $\text{parm\_fwdarg}[n_{\text{pth}}].\tilde{\mathbf{p}}_0 = \tilde{\mathbf{p}}_0[0]$ 
4:    $\text{parm\_fwdarg}[n_{\text{pth}}].\alpha_{\text{pth}} = \alpha_{\text{sph}}[n_{\text{pth}}N_xN_yN_z/N_{\text{pth}}]$ 
5:    $\text{parm\_fwdarg}[n_{\text{pth}}].\tilde{\mathbf{u}}'_{\text{pth}} = \tilde{\mathbf{u}}'_{\text{sph}}[n_{\text{pth}}N_rN_vN_f]$  {Pass addresses of arrays to each pthread}
6:    $\text{pthread\_create}(\&\text{pthreads}[n_{\text{pth}}], \text{NULL}, \text{fwd\_pthread}, (\text{void} *) (\text{parm\_fwdarg} + n_{\text{pth}}))$  {Call function fwd_pthread}
7: end for
8: for  $n_{\text{pth}} = 0$  to  $N_{\text{pth}} - 1$  do
9:   for  $n = 0$  to  $N_rN_vN_f$  do
10:     $\tilde{\mathbf{u}}_{\text{sph}}[n] += \tilde{\mathbf{u}}'_{\text{pth}}[n + n_{\text{pth}}N_rN_vN_f]$ 
11:   end for
12: end for

```


ALGORITHM VIII. Implementation of function fwd_pthread (on host).

Input: $n_{\text{pth}}, \tilde{\mathbf{p}}_0, \alpha_{\text{pth}}, \tilde{\mathbf{u}}'_{\text{pth}}$
Output: $\tilde{\mathbf{u}}'_{\text{pth}}$

```

1:  $\mathbf{C\_p0} \leftarrow \tilde{\mathbf{p}}_0$  {Copy from host to constant memory}
2: for  $n_x = 0$  to  $N_x/N_{\text{pth}} - 1$  do
3:    $x = (n_x + n_{\text{pth}}N_x/N_{\text{pth}})\Delta_s + x_{\min}$ 
4:   for  $n_y = 0$  to  $N_y - 1$  do
5:      $y = n_y\Delta_s + y_{\min}$ 
6:      $\mathbf{C\_apth} \leftarrow \alpha_{\text{pth}}[n_x][n_y][:]$  {Copy from host to constant memory}
7:      $\mathbf{K\_fwdspth} \langle \langle (N_v, N_r), (N_f, 1, 1) \rangle \rangle (x, y, \mathbf{D\_u'_{pth}})$ 
8:   end for
9: end for
10:  $\tilde{\mathbf{u}}'_{\text{pth}} \leftarrow \mathbf{D\_u'_{pth}}$  {Copy from global memory to host}
```

IV.A.3. Reconstruction algorithms

The GPU accelerated FBP algorithm was employed to reconstruct the object function sampled on a 3D Cartesian grid with spacing $\Delta_s = 0.14$ mm. The dimension of the reconstructed images $\hat{\alpha}_{\text{fbp}}$ was $210 \times 210 \times 440$.

We employed an iterative image reconstruction algorithm that sought to minimize a penalized least-squares (PLS) objective.^{45,50} Two versions of the reconstruction algorithm were developed that utilized the interpolation-based imaging model and the spherical-voxel-based imaging model, respectively. The two versions sought to solve the optimization problems by use of the linear conjugate gradient (CG) method^{51,52}

$$\hat{\alpha}_{\text{int}} = \arg \min_{\alpha_{\text{int}}} \|\mathbf{u} - \mathbf{H}_{\text{int}}\alpha_{\text{int}}\|^2 + \mu R(\alpha_{\text{int}}) \quad (29)$$

and

$$\hat{\alpha}_{\text{sph}} = \arg \min_{\alpha_{\text{sph}}} \|\tilde{\mathbf{u}} - \mathbf{H}_{\text{sph}}\alpha_{\text{sph}}\|^2 + \mu R(\alpha_{\text{sph}}), \quad (30)$$

respectively, where $R(\alpha)$ is a regularizing penalty term whose impact is controlled by the regularization parameter μ . The penalty term was employed only when processing the experimental data as described in Sec. IV.B. The reconstruction

ALGORITHM X. Implementation of $\alpha'_{\text{sph}} = \mathbf{H}_{\text{sph}}^{\dagger} \tilde{\mathbf{u}}$ (on host).

Input: $\tilde{\mathbf{u}}, \tilde{\mathbf{p}}_0$
Output: α'_{sph}

```

1: for  $n_{\text{pth}} = 0$  to  $N_{\text{pth}} - 1$  do
2:    $\text{parm\_bwdarg}[n_{\text{pth}}].n_{\text{pth}} = n_{\text{pth}}$ 
3:    $\text{parm\_bwdarg}[n_{\text{pth}}].\tilde{\mathbf{p}}_0 = \tilde{\mathbf{p}}_0[0]$ 
4:    $\text{parm\_bwdarg}[n_{\text{pth}}].\tilde{\mathbf{u}}_{\text{pth}} = \tilde{\mathbf{u}}[n_{\text{pth}}N_rN_vN_f/N_{\text{pth}}]$ 
5:    $\text{parm\_bwdarg}[n_{\text{pth}}].\alpha''_{\text{sph}} = \alpha''_{\text{sph}}[n_{\text{pth}}N_xN_yN_z]$ 
   {Pass addresses of arrays to each pthread}
6:    $\text{pthread\_create}(\&\text{pthreads}[n_{\text{pth}}], \text{NULL}, \text{bwd\_pthread},$ 
    $(\text{void}^*)(\text{parm\_bwdarg} + n_{\text{pth}}))$  {Call function bwd_pthread}
7: end for
8: for  $n_{\text{pth}} = 0$  to  $N_{\text{pth}} - 1$  do
9:   for  $n = 0$  to  $N_xN_yN_z$  do
10:     $\alpha'_{\text{sph}}[n] += \alpha''_{\text{sph}}[n + n_{\text{pth}}N_xN_yN_z]$ 
11:   end for
12: end for
```

ALGORITHM XI. Implementation of function bwd_pthread (on host).

Input: $n_{\text{pth}}, \tilde{\mathbf{p}}_0, \tilde{\mathbf{u}}_{\text{pth}}, \alpha''_{\text{pth}}$
Output: α''_{pth}

```

1:  $\mathbf{C\_p0} \leftarrow \tilde{\mathbf{p}}_0$  {Copy from host to constant memory}
2: for  $n_r = 0$  to  $N_r/N_{\text{pth}} - 1$  do
3:    $\theta^s = (n_r + n_{\text{pth}}N_r/N_{\text{pth}})\Delta_{\theta^s} + \theta_{\min}^s; z^s = R^s \cos \theta^s; r^s = R^s \sin \theta^s$ 
4:   for  $n_v = 0$  to  $N_v - 1$  do
5:      $\phi^s = n_v\Delta_{\phi^s} + \phi_{\min}^s; x^s = r^s \cos \phi^s; y^s = r^s \sin \phi^s$ 
6:      $\mathbf{C\_u_{pth}} \leftarrow \tilde{\mathbf{u}}_{\text{pth}}[n_r][n_v][:]$  {Copy from host to constant memory}
7:      $\mathbf{K\_bwdspth} \langle \langle (N_y, N_x), (N_z, 1, 1) \rangle \rangle (x^s, y^s, z^s, \mathbf{D\_u'_{pth}})$ 
8:   end for
9: end for
10:  $\alpha''_{\text{pth}} \leftarrow \mathbf{D\_u'_{pth}}$  {Copy from global memory to host}
```

algorithms required computation of one projection and one backprojection operation at each iteration. Hereafter, the two reconstruction algorithms will be referred to as PLS-Int and PLS-Sph algorithms, respectively. We set $\Delta_s = 0.14$ mm. Therefore, both the dimensions of $\hat{\alpha}_{\text{int}}$ and $\hat{\alpha}_{\text{sph}}$ were $210 \times 210 \times 440$.

ALGORITHM IX. Implementation of Kernel $\mathbf{K_fwdspth} \langle \langle (N_v, N_r), (N_f, 1, 1) \rangle \rangle$.

Input: $x, y, \mathbf{D_u'_{pth}}, \mathbf{C_apth}, \mathbf{C_p0}$
Output: $\mathbf{D_u'_{pth}}$

```

1:  $\theta^s = (\text{blockIdx.y})\Delta_{\theta^s} + \theta_{\min}^s; \phi^s = (\text{blockIdx.x})\Delta_{\phi^s} + \phi_{\min}^s; f = (\text{threadIdx.x})\Delta_f + f_{\min}$ 
2:  $z^s = R^s \cos \theta^s; x^s = R^s \sin \theta^s \cos \phi^s; y^s = R^s \sin \theta^s \sin \phi^s$  {Calculate locations of transducers}
3:  $\Sigma^r = 0; \Sigma^i = 0$  {Initiate the partial summation including the real and imaginary parts}
4: for  $n_z = 0$  to  $N_z - 1$  do
5:    $z = n_z\Delta_s + z_{\min}$ 
6:    $d = ((x - x^s)^2 + (y - y^s)^2 + (z - z^s)^2)^{1/2}$ 
7:    $\tilde{h}^r = \cos(2\pi f d / c_0) / (2\pi d); \tilde{h}^i = -\sin(2\pi f d / c_0) / (2\pi d)$  {Calculate SIR; Example here assumes point-like transducers}
8:    $\Sigma^r += \mathbf{C\_apth}[n_z](\tilde{h}^r \mathbf{C\_p0}[\text{threadIdx.x}].r - \tilde{h}^i \mathbf{C\_p0}[\text{threadIdx.x}].i)$ 
9:    $\Sigma^i += \mathbf{C\_apth}[n_z](\tilde{h}^i \mathbf{C\_p0}[\text{threadIdx.x}].i + \tilde{h}^r \mathbf{C\_p0}[\text{threadIdx.x}].r)$ 
10: end for
11:  $\mathbf{D\_u'_{pth}}[\text{blockIdx.y}][\text{blockIdx.x}][\text{threadIdx.x}].r += \Sigma^r$ 
12:  $\mathbf{D\_u'_{pth}}[\text{blockIdx.y}][\text{blockIdx.x}][\text{threadIdx.x}].i += \Sigma^i$ 
```

ALGORITHM XII. Implementation of Kernel $K_bwdsph \langle \langle (N_y, N_x), (N_z, 1, 1) \rangle \rangle$.

Input: $x^s, y^s, z^s, D_a''_{pth}, C_u_{pth}, C_p0$

Output: $D_a''_{pth}$

```

1:  $x = (\text{blockIdx.y})\Delta_s + x_{\min}; \quad y = (\text{blockIdx.x})\Delta_s + y_{\min}; \quad z = (\text{threadIdx.x})\Delta_s + z_{\min}$ 
2:  $d = ((x - x^s)^2 + (y - y^s)^2 + (z - z^s)^2)^{1/2}; \quad \Sigma = 0$  {Initiate the partial summation}
3: for  $n_f = 0$  to  $N_f - 1$  do
4:    $f = n_f\Delta_f + f_{\min}$ 
5:    $\tilde{h}^r = \cos(2\pi f d / c_0) / (2\pi d); \quad \tilde{h}^i = -\sin(2\pi f d / c_0) / (2\pi d)$  {Calculate SIR; Example here assumes point-like transducers}
6:    $\Sigma += C\_u_{pth}[n_f].r(\tilde{h}^r C\_p0[n_f].r - \tilde{h}^i C\_p0[n_f].i) + C\_u_{pth}[n_f].i(\tilde{h}^i C\_p0[n_f].r + \tilde{h}^r C\_p0[n_f].i)$ 
7: end for
8:  $D\_a''_{pth}[\text{blockIdx.y}][\text{blockIdx.x}][\text{threadIdx.x}] += \Sigma$ 

```

IV.A.4. Performance assessment

We compared the computational times of 3D image reconstruction corresponding to the GPU- and CPU-based implementations. The CPU-based implementations of the PLS-Int and PLS-Sph algorithms take several days to complete a single iteration even for the “32 × 45”-data. Therefore, we only recorded the computational time for the CPU-based implementations to complete a single iteration when the data vector contained a single transducer. We assumed that the computational times were linearly proportional to the number of transducers in the data sets because the CPU-based implementations are sequential.

The GPU-based implementations employed the single-precision floating-point format rather than the conventional double-precision utilized by CPU-based implementations. In order to quantify how the single-precision floating-point format would degrade the image accuracy, we calculated the root-mean-square-error (RMSE) between the reconstructed image and the phantom defined by

$$\text{RMSE} = \sqrt{\frac{1}{N}(\hat{\alpha} - \alpha)^T(\hat{\alpha} - \alpha)}, \quad (31)$$

where α and $\hat{\alpha}$ are the samples of the phantom and the coefficients of the reconstructed images, respectively.

IV.A.5. Hardware specifications

All implementations were tested on the platform consisted of dual quad-core Intel(R) Xeon (R) CPUs with a clock speed 2.40 GHz. The GPU-based implementations of the FBP and the PLS-Int algorithms were tested on a single Tesla C2050 GPU, while the PLS-Sph algorithm was tested on 8 Tesla C1060 GPUs.

IV.B. Experimental studies

The FBP, PLS-Int, and PLS-Sph algorithms were investigated by use of an existing data set corresponding to a live mouse.^{22,48} The scanning geometry and dimensions were the same as those employed in the computer-simulation studies except that only 64 rings were uniformly distributed over the polar angle ranging from 14° to 83°. The transducers were of size 2 × 2 mm². The raw data were acquired at 180 tomographic views, which are referred to as “full data.”

We undersampled the “full data” uniformly over the tomographic views, constructing a subset containing 45 tomographic views. The subset will be referred to as “quarter data.”

Unlike in the idealized computer-simulation studies, the transducer response has to be compensated for when processing the experimental data. When implementing the FBP algorithm, the EIR was compensated for by a direct Fourier deconvolution, expressed in temporal frequency domain as³

$$\tilde{p}(\mathbf{r}^s, f) = \frac{\tilde{u}(\mathbf{r}^s, f)}{\tilde{h}^e(f)} \tilde{W}(f), \quad (32)$$

where $\tilde{W}(f)$ is a window function for noise suppression. In this study, we adopted the Hann window function defined as

$$\tilde{W}(f) = \frac{1}{2} \left[1 - \cos \left(\pi \frac{f_c - f}{f_c} \right) \right], \quad (33)$$

where the cutoff frequency was chosen as $f_c = 5$ MHz. When applying iterative image reconstruction algorithms, the transducer effects were implicitly compensated for during iteration by employing imaging models that incorporates the transducer characteristics.^{14,22} We incorporated the EIR into the interpolation-based imaging model while incorporating both the EIR and the SIR into the spherical-voxel-based imaging model.

For both PLS-Int and PLS-Sph algorithms, we employed a quadratic smoothness penalty to mitigate measurement noise⁵⁰

$$R(\alpha) = \sum_{n=0}^{N-1} ([\alpha]_n - [\alpha]_{n_x})^2 + ([\alpha]_n - [\alpha]_{n_y})^2 + ([\alpha]_n - [\alpha]_{n_z})^2, \quad (34)$$

where n_x , n_y , and n_z were the indices of the neighboring voxels before the n th voxel along the three Cartesian axes, respectively.

V. RESULTS

V.A. Computational efficiency

As shown in Table I, the GPU-based implementations took less than 0.1%, 0.8%, and 0.4% of the computational times required by corresponding CPU-based implementations for the FBP, the PLS-Int, and the PLS-Sph algorithms, respectively. The relative computational times for the GPU-based

TABLE I. Computational times of the 3D image reconstructions by use of the CPU- and GPU-based implementations.

Data sets	FBP [s]		PLS-Int [min/iteration]		PLS-Sph [min/iteration]	
	CPU	GPU	CPU	GPU	CPU	GPU
"32 × 45"	6189	6	2448	20	7961	22
"64 × 45"	12 975	12	4896	35	15 923	43
"64 × 90"	26 190	23	9792	68	31 845	86
"128 × 90"	53 441	46	...	...	...	...
"Quarter data"	12 975	12	4896	35	19 776	78
"Full data"	53 441	46	19 968	137	79 177	313

implementations are nearly linearly proportional to the amount of data. Note that the "64 × 90"-data and the "quarter data" are of the same size. However, the computational times of the "quarter data" are more than 1.8 times those of the "64 × 90"-data. This is because the calculation of the SIR increases the computational complexity of the reconstruction algorithm.

V.B. Computational accuracy

Images reconstructed by use of the CPU- and GPU-based implementations of the FBP algorithm are almost identical. From the "128 × 90"-data, in which case, transducers were densely distributed over the measurement surface, both implementations reconstructed accurate images, as shown in Figs. 2(b) and 2(c). The profiles along the three arrows in Fig. 2 are plotted in Fig. 5(a), suggesting a nearly exact reconstruction. As expected, when the amount of measurement data are reduced, the reconstructed images contain more artifacts as shown in Fig. 3. However, the images reconstructed by use of GPU- and CPU-based implementations remain indistinguishable. The plots of the RMSE versus the amount of measurement data employed in Fig. 6 overlap, also suggesting the single-precision floating-point format employed by the GPU-based implementation has little impact on the computational accuracy.

The GPU-based implementations of the PLS-Int and PLS-Sph algorithms both reconstructed accurate images as displayed in Fig. 4. As expected, the images reconstructed by use of both iterative algorithms contain fewer artifacts than those reconstructed by use of the FBP algorithm from the same

amount of data. Unlike the images reconstructed by use of the FBP algorithm from the "64 × 90"-data [Figs. 3(a) or 3(d)], the images reconstructed by use of both iterative algorithms [Figs. 4(a) and 4(d)] appear to be identical to the numerical phantom. The profiles along the two arrows in Figs. 4(a) and 4(d) are plotted in Fig. 5(b), further confirming the computational accuracy of iterative image reconstruction algorithms. The plots of the RMSE versus the amount of measurement data employed in Fig. 6 suggest the iterative image reconstruction algorithms in general outperform the FBP algorithm from the same amount of data.

V.C. Experimental results

The maximum intensity projection (MIP) of the 3D mouse images reconstructed by use of the GPU-based implementations reveal the mouse body vasculature as shown in Fig. 7.

FIG. 2. Slices corresponding to the plane $y = 0$ of (a) the phantom and the images reconstructed by use of (b) the CPU-based and (c) the GPU-based implementations of the FBP algorithm from the "128 × 90"-data.

FIG. 3. Slices corresponding to the plane $y = 0$ of the images reconstructed by use of the FBP algorithm with (a) the CPU-based implementation from the "64 × 90"-data, (b) the CPU-based implementation from the "64 × 45"-data, (c) the CPU-based implementation from the "32 × 45"-data, (d) the GPU-based implementation from the "64 × 90"-data, (e) the GPU-based implementation from the "64 × 45"-data, and (f) the GPU-based implementation from the "32 × 45"-data.

FIG. 4. Slices corresponding to the plane $y = 0$ of the images reconstructed by use of the GPU-based implementations of (a) the PLS-Int algorithm from the “ 64×90 ”-data, (b) the PLS-Int algorithm from the “ 64×45 ”-data, (c) the PLS-Int algorithm from the “ 32×45 ”-data, (d) the PLS-Sph algorithm from the “ 64×90 ”-data, (e) the PLS-Sph algorithm from the “ 64×45 ”-data, and (f) the PLS-Sph algorithm from the “ 32×45 ”-data.

FIG. 5. Profiles along the line $(x, y) = (-6.58, 0)$ mm of the images reconstructed by use of (a) the CPU- and GPU-based implementations of the FBP algorithm from the “ 128×90 ”-data, and (b) the GPU-based implementations of the PLS-Int and the PLS-Sph algorithms from the “ 64×90 ”-data.

FIG. 6. Plots of the RMSE against the amount of data by use of the FBP, the PLS-Int and the PLS-Sph algorithms.

FIG. 7. MIP renderings of the 3D images of the mouse body reconstructed by use of the GPU-based implementations of (a) the FBP algorithm from the “full data,” (b) the PLS-Int algorithm from the “full data” with $\mu = 1.0 \times 10^4$, (c) the PLS-Sph algorithm from the “full data” with $\mu = 1.0 \times 10^4$, (d) the FBP algorithm from the “quarter data,” (e) the PLS-Int algorithm from the “quarter data” with $\mu = 1.0 \times 10^3$, and (f) the PLS-Sph algorithm from the “quarter data” with $\mu = 1.0 \times 10^3$. The grayscale window is $[0, 12.0]$.

Images reconstructed by use of both the PLS-Int and the PLS-Sph algorithms appear to have cleaner background than the images reconstructed by use of the FBP algorithm from the same amount of data. All images reconstructed by iterative algorithms were obtained by 20-iterations starting with uniform zeros as the initial guess. The PLS-Int algorithm took approximately a half day and 2 days to process the “quarter data” and the “full data,” respectively. The PLS-Sph algorithm took approximately 1 day and 4 days to process the “quarter data” and the “full data,” respectively. Alternatively, if the CPU-based implementations were utilized, the PLS-Int algorithm would take an estimated 68 days and 277 days to process the “quarter data” and the “full data,” respectively. The PLS-Sph algorithm would take an estimated 275 days and 1100 days to process the “quarter data” and the “full data,” respectively.

VI. DISCUSSION AND CONCLUSION

In this study, we developed and investigated GPU-based implementations of the FBP algorithm and two pairs of projection/backprojection operators for 3D OAT. Our implementation of the FBP algorithm improved the computational efficiency over 1000 times compared to the CPU-based implementation. This work complements our earlier studies that demonstrated the feasibility of 3D iterative image reconstruction in practice.^{21,22}

Our current implementations of the iterative image reconstruction algorithms still require several days to process the densely sampled data set, which, however, can be further improved. First, the amount of measurement data required for accurate image reconstruction can be further reduced by developing advanced image reconstruction methods.^{13,15,22,53} Second, the number of iterations required can be reduced by developing fast-converging optimization algorithms.^{22,54,55}

The proposed parallelization strategies by use of GPUs are of general interest. The implementation of the FBP algorithm⁶ can be adapted to other analytic image reconstruction algorithms, including those described in Refs. 5, 7, 56–58. We demonstrated the feasibility of PLS algorithm that utilized the proposed GPU-based implementations of the projection/backprojection operators. By use of these implementations, many advanced image reconstruction algorithms may also be feasible in practice.²² Though we described our parallelization strategies for the projection/backprojection operators that utilized two discrete-to-discrete imaging models, these strategies can also be applied to other D-D imaging models.^{9,11,20,59} Therefore, the proposed algorithms will facilitate the further investigation and application of advanced image reconstruction algorithms in 3D OAT.

ACKNOWLEDGMENT

This research was supported in part by the National Institutes of Health (NIH) Award Nos. EB010049 and CA167446.

APPENDIX: DISCRETIZATION OF SRT

1. Derivation of Eq. (12)

The integrated data function $g(\mathbf{r}^s, t)$ in Eq. (2), evaluated at the q th transducer and the k th time instance, can be expressed as

$$g(\mathbf{r}_q^s, t)|_{t=k\Delta_t} = \int_{|\mathbf{r}_q^s - \mathbf{r}|=kc_0\Delta_t} d\mathbf{r} A(\mathbf{r}), \quad (\text{A1})$$

where $\mathbf{r}_q^s$ denotes the location of the q th point-like transducer. We defined a local coordinate system, distinguished by a superscript “tr,” centered at the q th transducer with the z^{tr} -axis pointing to the origin of the global coordinate system as shown in Fig. 1(b). Assuming the object function $A(\mathbf{r})$ is compactly supported in a sphere of radius R , the integral surface is symmetric about the z^{tr} -axis. Thus, the orientations of the x^{tr} - and y^{tr} -axes can be arbitrary within the $z^{\text{tr}} = 0$ plane. Representing the right-hand side of Eq. (A1) in the local spherical coordinate system, one obtains

$$g(\mathbf{r}_q^s, t)|_{t=k\Delta_t} = (kc_0\Delta_t)^2 \int_0^{\theta_{\text{max}}^{\text{tr}}} d\theta^{\text{tr}} \sin \theta^{\text{tr}} \times \int_0^{2\pi} d\phi^{\text{tr}} A(kc_0\Delta_t, \theta^{\text{tr}}, \phi^{\text{tr}}), \quad (\text{A2})$$

where $\theta_{\text{max}}^{\text{tr}}$ is half of the apex angle of the cone that corresponds to the intersectional spherical cap as shown in Fig. 1(b). The polar angle θ^{tr} and the azimuth angle ϕ^{tr} were discretized with intervals $\Delta\theta^{\text{tr}}$ and $\Delta\phi^{\text{tr}}$ that satisfied

$$kc_0\Delta_t\Delta\theta^{\text{tr}} = kc_0\Delta_t\sin\theta^{\text{tr}}\Delta\phi^{\text{tr}} = \Delta_s. \quad (\text{A3})$$

Therefore, Eq. (A2) can be approximated by

$$g(\mathbf{r}_q^s, t)|_{t=k\Delta_t} \approx \Delta_s^2 \sum_{i=0}^{N_i-1} \sum_{j=0}^{N_j-1} A(kc_0\Delta_t, \theta_i^{\text{tr}}, \phi_j^{\text{tr}}), \quad (\text{A4})$$

where $N_i = \lfloor \theta_{\text{max}}^{\text{tr}} / \Delta\theta^{\text{tr}} \rfloor$, $N_j = \lfloor 2\pi / \Delta\phi^{\text{tr}} \rfloor$, $\theta_i^{\text{tr}} = i\Delta\theta^{\text{tr}}$, and $\phi_j^{\text{tr}} = j\Delta\phi^{\text{tr}}$. We denoted by $\mathbf{r}_{k,i,j}$ the location in the global coordinate system corresponding to the location vector $(kc_0\Delta_t, \theta_i^{\text{tr}}, \phi_j^{\text{tr}})$ in the local coordinate system in Eq. (A4). On substitution from the finite-dimensional representation Eq. (6) into Eq. (A4) with α and $\psi_n(\mathbf{r})$ defined by Eqs. (9) and (10), respectively, we obtained

$$g(\mathbf{r}_q^s, t)|_{t=k\Delta_t} \approx \Delta_s^2 \sum_{n=0}^{N-1} [\alpha_{\text{int}}]_n \sum_{i=0}^{N_i-1} \sum_{j=0}^{N_j-1} \psi_n^{\text{int}}(\mathbf{r}_{k,i,j}) \equiv [\mathbf{g}]_{qK+k}. \quad (\text{A5})$$

^{a)}Electronic mail: anastasio@seas.wustl.edu

¹A. A. Oraevsky and A. A. Karabutov, “Optoacoustic tomography,” in *Biomedical Photonics Handbook*, edited by T. Vo-Dinh (CRC, Boca Raton, FL, 2003), Chap. 34.

²L. V. Wang, “Tutorial on photoacoustic microscopy and computed tomography,” *IEEE J. Sel. Top. Quantum Electron.* **14**, 171–179 (2008).

³R. Kruger, D. Reinecke, and G. Kruger, “Thermoacoustic computed tomography: Technical considerations,” *Med. Phys.* **26**, 1832–1837 (1999).

- ⁴B. T. Cox, S. R. Arridge, K. P. Köstli, and P. C. Beard, "Two-dimensional quantitative photoacoustic image reconstruction of absorption distributions in scattering media by use of a simple iterative method," *Appl. Opt.* **45**, 1866–1875 (2006).
- ⁵L. A. Kunyansky, "Explicit inversion formulae for the spherical mean Radon transform," *Inverse Probl.* **23**, 373–383 (2007).
- ⁶D. Finch, S. Patch, and Rakesh, "Determining a function from its mean values over a family of spheres," *SIAM J. Math. Anal.* **35**, 1213–1240 (2004).
- ⁷M. Xu and L. V. Wang, "Universal back-projection algorithm for photoacoustic computed tomography," *Phys. Rev. E* **71**, 016706 (2005).
- ⁸Y. Xu, D. Feng, and L. V. Wang, "Exact frequency-domain reconstruction for thermoacoustic tomography: I. Planar geometry," *IEEE Trans. Med. Imaging* **21**, 823–828 (2002).
- ⁹G. Paltouf, J. A. Viator, S. A. Prahl, and S. L. Jacques, "Iterative reconstruction algorithm for optoacoustic imaging," *J. Acoust. Soc. Am.* **112**, 1536–1544 (2002).
- ¹⁰Z. Yuan and H. Jiang, "Three-dimensional finite-element-based photoacoustic tomography: Reconstruction algorithm and simulations," *Med. Phys.* **34**, 538–546 (2007).
- ¹¹P. Ephrat, L. Keenliside, A. Seabrook, F. S. Prato, and J. J. L. Carson, "Three-dimensional photoacoustic imaging by sparse-array detection and iterative image reconstruction," *J. Biomed. Opt.* **13**, 054052 (2008).
- ¹²J. Zhang, M. Anastasio, P. La Riviere, and L. Wang, "Effects of different imaging models on least-squares image reconstruction accuracy in photoacoustic tomography," *IEEE Trans. Med. Imaging* **28**, 1781–1790 (2009).
- ¹³J. Provost and F. Lesage, "The application of compressed sensing for photo-acoustic tomography," *IEEE Trans. Med. Imaging* **28**, 585–594 (2009).
- ¹⁴K. Wang, S. A. Ermilov, R. Su, H.-P. Brecht, A. A. Oraevsky, and M. A. Anastasio, "An imaging model incorporating ultrasonic transducer properties for three-dimensional optoacoustic tomography," *IEEE Trans. Med. Imaging* **30**, 203–214 (2011).
- ¹⁵Z. Guo, C. Li, L. Song, and L. V. Wang, "Compressed sensing in photoacoustic tomography *in vivo*," *J. Biomed. Opt.* **15**, 021311 (2010).
- ¹⁶C. Huang, A. A. Oraevsky, and M. A. Anastasio, "Investigation of limited-view image reconstruction in optoacoustic tomography employing *a priori* structural information," *Proc. SPIE* **7800**, 780004 (2010).
- ¹⁷Z. Xu, C. Li, and L. V. Wang, "Photoacoustic tomography of water in phantoms and tissue," *J. Biomed. Opt.* **15**, 036019 (2010).
- ¹⁸Z. Xu, Q. Zhu, and L. V. Wang, "In vivo photoacoustic tomography of mouse cerebral edema induced by cold injury," *J. Biomed. Opt.* **16**, 066020 (2011).
- ¹⁹A. Buehler, A. Rosenthal, T. Jetzfellner, A. Dima, D. Razansky, and V. Ntziachristos, "Model-based optoacoustic inversions with incomplete projection data," *Med. Phys.* **38**, 1694–1704 (2011).
- ²⁰S. Bu, Z. Liu, T. Shiina, K. Kondo, M. Yamakawa, K. Fukutani, Y. Someda, and Y. Asao, "Model-based reconstruction integrated with fluence compensation for photoacoustic tomography," *IEEE Trans. Biomed. Eng.* **59**, 1354–1363 (2012).
- ²¹K. Wang, R. Su, A. A. Oraevsky, and M. A. Anastasio, "Investigation of iterative image reconstruction in optoacoustic tomography," *Proc. SPIE* **8223**, 82231Y (2012).
- ²²K. Wang, R. Su, A. A. Oraevsky, and M. A. Anastasio, "Investigation of iterative image reconstruction in three-dimensional optoacoustic tomography," *Phys. Med. Biol.* **57**, 5399 (2012).
- ²³C. Huang, L. Nie, R. W. Schoonover, Z. Guo, C. O. Schirra, M. A. Anastasio, and L. V. Wang, "Aberration correction for transcranial photoacoustic tomography of primates employing adjunct image data," *J. Biomed. Opt.* **17**, 066016 (2012).
- ²⁴X. Dean-Ben, A. Buehler, V. Ntziachristos, and D. Razansky, "Accurate model-based reconstruction algorithm for three-dimensional optoacoustic tomography," *IEEE Trans. Med. Imaging* **31**, 1922–1928 (2012).
- ²⁵C. Huang, L. Nie, R. W. Schoonover, L. V. Wang, and M. A. Anastasio, "Photoacoustic computed tomography correcting for heterogeneity and attenuation," *J. Biomed. Opt.* **17**, 061211 (2012).
- ²⁶B. Wang, L. Xiang, M. S. Jiang, J. Yang, Q. Zhang, P. R. Carney, and H. Jiang, "Photoacoustic tomography system for noninvasive real-time three-dimensional imaging of epilepsy," *Biomed. Opt. Express* **3**, 1427–1432 (2012).
- ²⁷A. Buehler, X. L. Deán-Ben, J. Claussen, V. Ntziachristos, and D. Razansky, "Three-dimensional optoacoustic tomography at video rate," *Opt. Express* **20**, 22712–22719 (2012).
- ²⁸E. Lindholm, J. Nickolls, S. Oberman, and J. Montrym, "Nvidia tesla: A unified graphics and computing architecture," *IEEE MICRO* **28**, 39–55 (2008).
- ²⁹NVIDIA CUDA Programming Guide 2.0 NVIDIA, 2008.
- ³⁰X. Zhao, J.-J. Hu, and P. Zhang, "GPU-based 3D cone-beam CT image reconstruction for large data volume," *J. Biomed. Imaging* **2009**(8), 1–8 (2009).
- ³¹Y. Okitsu, F. Ino, and K. Hagihara, "High-performance cone beam reconstruction using CUDA compatible GPUs," *Parallel Comput.* **36**, 129–141 (2010).
- ³²C.-Y. Chou, Y.-Y. Chuo, Y. Hung, and W. Wang, "A fast forward projection using multithreads for multirays on GPUs in medical image reconstruction," *Med. Phys.* **38**, 4052–4065 (2011).
- ³³S. Stone, J. Haldar, S. Tsao, W.-M. W. Hwu, B. Sutton, and Z.-P. Liang, "Accelerating advanced MRI reconstructions on GPUs," *J. Parallel Distrib. Comput.* **68**, 1307–1318 (2008).
- ³⁴B. E. Treeby and B. T. Cox, "k-wave: plbisc-matlab toolbox for the simulation and reconstruction of photoacoustic wave fields," *J. Biomed. Opt.* **15**, 021314 (2010).
- ³⁵H. Barrett and K. Myers, *Foundations of Image Science*, Wiley Series in Pure and Applied Optics (John Wiley & Son, Hoboken, NJ, 2004).
- ³⁶L. V. Wang and H.-I. Wu, *Biomedical Optics, Principles and Imaging* (Wiley, Hoboken, NJ, 2007).
- ³⁷M. Xu and L. V. Wang, "Photoacoustic imaging in biomedicine," *Rev. Sci. Instrum.* **77**, 041101–041122 (2006).
- ³⁸K. Wang, and M. A. Anastasio, "Photoacoustic and thermoacoustic tomography: Image formation principles," in *Handbook of Mathematical Methods in Imaging*, edited by O. Scherzer (Springer, New York, NY, 2011), Chap. 18.
- ³⁹A. Rosenthal, V. Ntziachristos, and D. Razansky, "Optoacoustic methods for frequency calibration of ultrasonic sensors," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **58**, 316–326 (2011).
- ⁴⁰A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging* (IEEE, New York, NY, 1988).
- ⁴¹M. A. Anastasio, J. Zhang, X. Pan, Y. Zou, G. Keng, and L. V. Wang, "Half-time image reconstruction in thermoacoustic tomography," *IEEE Trans. Med. Imaging* **24**, 199–210 (2005).
- ⁴²M. Anastasio, J. Zhang, E. Sidky, Y. Zou, D. Xia, and X. Pan, "Feasibility of half-data image reconstruction in 3-D reflectivity tomography with a spherical aperture," *IEEE Trans. Med. Imaging* **24**, 1100–1112 (2005).
- ⁴³J. F. Claerbout, *Earth Sounding Analysis: Processing Versus Inversion* (Blackwell Scientific, Cambridge, MA, 1992).
- ⁴⁴T. D. Khokhlova, I. M. Pelivanov, V. V. Kozhushko, A. N. Zharinov, V. S. Solomatin, and A. A. Karabutov, "Optoacoustic imaging of absorbing objects in a turbid medium: Ultimate sensitivity and application to breast cancer diagnostics," *Appl. Opt.* **46**, 262–272 (2007).
- ⁴⁵J. N. Aarsvold, *Emission Tomography: The Fundamentals of PET and SPECT* (Elsevier Academic, San Diego, CA, 2004).
- ⁴⁶G. Zeng and G. Gullberg, "Unmatched projector/backprojector pairs in an iterative reconstruction algorithm," *IEEE Trans. Med. Imaging* **19**, 548–555 (2000).
- ⁴⁷K. W. Morton and D. F. Mayers, *Numerical Solution of Partial Differential Equations: An Introduction* (Cambridge University Press, New York, NY, 2005).
- ⁴⁸H.-P. Brecht, R. Su, M. Fronheiser, S. A. Ermilov, A. Conjusteau, and A. A. Oraevsky, "Whole-body three-dimensional optoacoustic tomography system for small animals," *J. Biomed. Opt.* **14**, 064007 (2009).
- ⁴⁹M. A. Anastasio, J. Zhang, D. Modgil, and P. La Riviere, "Application of inverse source concepts to photoacoustic tomography," *Inverse Probl.* **23**, S21–S35 (2007).
- ⁵⁰J. A. Fessler, "Penalized weighted least-squares reconstruction for positron emission tomography," *IEEE Trans. Med. Imaging* **13**, 290–300 (1994).
- ⁵¹J. R. Shewchuk, "An introduction to the conjugate gradient method without the agonizing pain," Technical Report No. CMU-CS-94-125 (Carnegie Mellon University, Pittsburgh, PA, 1994).
- ⁵²J. Fessler and S. Booth, "Conjugate-gradient preconditioning methods for shift-variant PET image reconstruction," *IEEE Trans. Image Process.* **8**, 688–699 (1999).
- ⁵³J. Meng, L. V. Wang, L. Ying, D. Liang, and L. Song, "Compressed-sensing photoacoustic computed tomography *in vivo* with partially known support," *Opt. Express* **20**, 16510–16523 (2012).

- ⁵⁴A. Beck and M. Teboulle, "Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems," *IEEE Trans. Image Process.* **18**, 2419–2434 (2009).
- ⁵⁵S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.* **3**, 1–122 (2011).
- ⁵⁶M. Xu and L. V. Wang, "Time-domain reconstruction for thermoacoustic tomography in a spherical geometry," *IEEE Trans. Med. Imaging* **21**, 814–822 (2002).
- ⁵⁷D. Finch, M. Haltmeier, and Rakesh, "Inversion of spherical means and the wave equation in even dimensions," *SIAM J. Appl. Math.* **68**, 392–412 (2007).
- ⁵⁸P. Elbau, O. Scherzer, and R. Schulze, "Reconstruction formulas for photoacoustic sectional imaging," *Inverse Probl.* **28**, 045004 (2012).
- ⁵⁹A. Rosenthal, D. Razansky, and V. Ntziachristos, "Fast semi-analytical model-based acoustic inversion for quantitative optoacoustic tomography," *IEEE Trans. Med. Imaging* **29**, 1275–1285 (2010).