Journal of Biomedical Optics

SPIEDigitalLibrary.org/jbo

Highly parallel Monte-Carlo simulations of the acousto-optic effect in heterogeneous turbid media

Samuel Powell
Terence S. Leung

Highly parallel Monte-Carlo simulations of the acousto-optic effect in heterogeneous turbid media

Samuel Powell and Terence S. Leung

University College London, Department of Medical Physics and Bioengineering, Malet Place Engineering Building, London, WC1E 6BT, United Kingdom

Abstract. The development of a highly parallel simulation of the acousto-optic effect is detailed. The simulation supports optically heterogeneous simulation domains under insonification by arbitrary monochromatic ultrasound fields. An adjoint method for acousto-optics is proposed to permit point-source/point-detector simulations. The flexibility and efficiency of this simulation code is demonstrated in the development of spatial absorption sensitivity maps which are in broad agreement with current experimental investigations. The simulation code has the potential to provide guidance in the feasibility and optimization of future studies of the acousto-optic technique, and its speed may permit its use as part of an iterative inversion model. © 2012 Society of Photo-Optical Instrumentation Engineers (SPIE). [DOI: 10.1117/1.JBO.17.4.045002]

Keywords: ultrasound modulated optical tomography; acousto-optics; Monte-Carlo; graphics processing unit; compute unified device architecture.

Paper 11647 received Nov. 6, 2011; revised manuscript received Jan. 27, 2012; accepted for publication Jan. 30, 2012; published online Apr. 6, 2012.

1 Introduction

Acousto-optics (AO), or ultrasound-modulated optical tomography, is a hybrid technique which combines the optical contrast of biological tissues at visible or near-infrared wavelengths with the spatial resolution of focused ultrasound fields to provide highly localized measurements of the optical properties of biological tissues. Such measurements are of clinical significance since they may indicate underlying pathologies, e.g., increased vascularization due to tumor angiogenesis, or serve as an indicator of the progress of treatments which themselves alter the optical properties of tissues, e.g., thermal necrosis by therapeutic high-intensity focused ultrasound (HIFU).¹

1.1 Acousto-Optic Modeling

Leutz and Maret² first developed a model of the AO effect in turbid media based upon the approach of diffusing wave spectroscopy (DWS)^{3,4} with the addition of sinusoidally oscillating optical scatterers. Kempe et al.⁵ extended the analysis of Leutz and Maret to consider a narrow beam of ultrasound and compared this model to experimental data. Mahan et al.⁶ proposed the acoustic tagging process to be that of Brillouin scattering, proceeding to develop an expression of the intensity of the modulated signal and the signal-to-noise ratio in the case of an ultrasound field focused into a small volume. Theoretical understanding of the technique was advanced by Wang,⁷ who jointly modeled the two coherent mechanisms of interaction previously investigated: the acoustically induced displacement of optical scatterers and the change in the index of refraction due to compression and rarefaction. Theoretical research continued with the extension of Wang's model to the case of anisotropically scattering media⁸ and to account for strong

correlations in the interaction at successive scattering sites.⁹ Blonigen et al.¹⁰ further investigated the nature of the phase shifts applied to ultrasound modulated light in the context of a photorefractive crystal detection regime. The inherent limitations of theory based upon the DWS approach were overcome by Sakadžić and Wang through the development of a correlation transfer equation (CTE)¹¹ which provides a treatment of the acoustic continuous-wave (CW) AO problem in the style of a transport equation. A diffusion approximation (DA) is presented in a subsequent work,¹² and later the CTE is reformulated to account for an ultrasonic pulse.¹³ Recently Bratchenia et al.¹⁴ developed sensitivity maps of the acousto-optic technique based upon a pair of coupled diffusion equations accounting for the unmodulated and modulated distributions of light. Throughout the initial theoretical investigations of the technique, a number of CPU-based Monte-Carlo (MC) models were demonstrated to support the analytical solutions under development.^{8,9,11–13,15}

1.2 GPU-Accelerated Monte-Carlo Technique

In recent years, considerable advances have been made in executing MC simulations on the highly parallel architecture of graphics processing units (GPUs). The first demonstration by Allerstam et al.¹⁶ considered the problem of light transport in an homogeneous semi-infinite medium; increases in speed over a CPU implementation of up to three orders of magnitude were achieved. Fang and Boas¹⁷ subsequently developed a general purpose code to perform MC simulations in heterogeneous voxelized geometries. Under specific circumstances speed increases over a CPU implementation of up to $\times 350$ were noted.

We recently demonstrated the acceleration of a MC AO simulation,¹⁸ achieving performance improvements of up to $\times 100$ over a CPU implementation equivalent to that presented in Ref. 15. We also demonstrated that our GPU-based

Address all correspondence to: Samuel Powell, University College London, Department of Medical Physics and Bioengineering, Malet Place Engineering Building, London, WC1E 6BT, United Kingdom. Tel: +44 20 7679 0273; E-mail: spowell@medphys.ucl.ac.uk.

AO simulations involved an order of magnitude of extra computational effort over a pure optical simulation.

1.3 Aims

Given the demonstrable improvements in efficiency which can be achieved using a GPU-based parallel implementation, our aim is the development of a robust, efficient and flexible simulation tool that has the potential to provide significant guidance in the application and optimization of the AO technique as both an imaging and a sensing modality for a broad range of tissue geometries and detection mechanisms. A fast forward model may also serve as the basis of future developments in inversion techniques for AO.

While our previous implementation¹⁸ was limited to homogeneous semi-infinite slab geometries with plane-wave ultrasound, here we develop a code capable of simulating the AO effect in optically heterogeneous domains, with arbitrary (monochromatic) acoustic field distributions. Previous CPU-based simulations reported in the literature^{11,13} have implemented a similar level of flexibility by the use of a voxelized simulation domain for the spatial distribution of the optical and acoustic parameters. Here we will demonstrate the implementation of optical heterogeneity via a mesh-based system which avoids the deleterious effects of a voxelized geometry,¹⁹ and exploits the computational power of the GPU to avoid approximations in the treatment of the AO phase shifts demonstrated in previous work.¹¹

It is our intention to release the described simulation code publicly in the near future; if you wish to be contacted upon its release, please contact the corresponding author by E-mail.

1.4 Overview

We begin in Sec. 2 with an overview of the relevant theory of AO, and the variance-reduction technique employed for point-source, point-detector (PSPD) simulations. The implementation of the simulation program and its post processing techniques are detailed in Sec. 3, following a brief overview of the pertinent considerations of parallel programming for GPU architectures. In Sec. 4 we demonstrate the validation of the simulation program, consider its performance relative to other simulation codes, and demonstrate one application in the calculation of explicit spatial sensitivity maps corresponding to experimental studies currently being undertaken. This work closes in Sec. 5 with a discussion of the limitations and possibilities for future development of the simulation code.

2 Theory

2.1 Acousto-Optic Interaction within Turbid Media

Three modes of AO interaction have been proposed in the literature: displacement of optical scatterers in a turbid medium by the acoustic field,² perturbation of the refractive index of a medium due to mechanical strain induced by the acoustic field,⁷ and perturbation of the optical properties of the medium due to compression and rarefaction of the medium by the acoustic field.⁶ The latter mechanism could in principle be detected with incoherent light, but this effect has been too weak to be detected in practice (without the use of fluorescence). We consider only the two former mechanisms, which require the use of coherent light.

2.1.1 Variation of the index of refraction

An acoustic wave propagating through a medium generates a strain which alters the local permittivity of the medium. The adiabatic piezo-optic coefficient $\eta = \delta n / \delta p$ ($\approx 1.466 \times 10^{-10}$ Pa⁻¹ in water) relates the change in the index of refraction n to the acoustic pressure p .²⁰ Consider a plane electromagnetic wave propagating between two points $\vec{r}_a$ and $\vec{r}_b$ with wavenumber $k_0 = \omega_o / c_0$, where ω_o is the angular frequency and c_0 the speed of light *in vacuo*. The incremental phase angle accrued on this path over that of an optically homogeneous medium of refractive index n_0 is given by

$$\phi_n(\vec{r}_a, \vec{r}_b, t) = k_0 n_0 \eta \int_{\vec{r}_a}^{\vec{r}_b} p(\vec{r}, t) d\vec{r}. \quad (1)$$

Under insonification by a monochromatic acoustic field of angular frequency ω_a , local pressure amplitude P_0 , and local phase offset relative to an arbitrary point $\phi_a(\vec{r})$ we express the complex time harmonic phase increment

$$\begin{aligned} \tilde{\phi}_n(\vec{r}_a, \vec{r}_b, t) &= k_0 n_0 \eta \int_{\vec{r}_a}^{\vec{r}_b} P_0(\vec{r}) \\ &\times \exp[j(\phi_a(\vec{r}) - \omega_a t)] d\vec{r}. \end{aligned} \quad (2)$$

2.1.2 Displacement of optical scatterers

An acoustic wave propagating through a turbid medium displaces the endogenous optical scatterers from their rest positions. The acoustic particle displacement at an arbitrary point in space and time is

$$\vec{\xi}_a(\vec{r}, t) = p(\vec{r}, t) \hat{k}_a / (\omega_a \rho v_a), \quad (3)$$

where $\hat{k}_a = \vec{k}_a / k_a$ is the unit acoustic propagation direction derived from the acoustic wave-vector k_a and wavenumber k_a , ρ is the density of the medium and v_a the speed of sound in the medium. We express the complex valued acoustic particle displacement for a monochromatic acoustic field

$$\tilde{\xi}_s(\vec{r}, t) = P_0(\vec{r}) \hat{k}_a / (\omega_a \rho v_a) \exp[j(\phi_a(\vec{r}) - \omega_a t)], \quad (4)$$

and assume that the displacement of a given scatterer follows that of the background medium with a particular phase offset ϕ_s and scaled amplitude S ,¹⁰

$$\begin{aligned} \tilde{\xi}_s(\vec{r}, t) &= S P_0(\vec{r}) \hat{k}_a / (\omega_a \rho v_a) \\ &\times \exp[j(\phi_a(\vec{r}) - \omega_a t + \phi_s)]. \end{aligned} \quad (5)$$

Assuming two scattering sites at $\vec{r}_a$ and $\vec{r}_b$ are separated such that they are in the far-field zone of one another, the phase increment accrued by an electromagnetic wave will be the projection of the optical wavevector $\vec{k}_0$ onto the overall scatterer displacement,

$$\tilde{\phi}_a(\vec{r}_a, \vec{r}_b, t) = n_0 \vec{k}_0 \cdot [\tilde{\xi}_s(\vec{r}_b, t) - \tilde{\xi}_s(\vec{r}_a, t)]. \quad (6)$$

2.1.3 Acousto-optic interaction in a multiple scattering regime

We may think of the electric field at a point within, or on the boundary of a coherently illuminated turbid medium to consist of the summation of an unscattered field resulting from the incident radiation and numerous partial fields which result from the multiple scattering process,

$$\tilde{E}(\vec{r}, t) = \tilde{E}_i(\vec{r}, t) + \sum_m \tilde{E}_m(\vec{r}, t), \quad (7)$$

where E_i denotes the unscattered field due to the incident radiation, and E_m the partial multiply scattered fields, each representing a contribution from a scattering process of some order. Direct calculation of the electric field by this approach is impractical for any realistic biological system. Transport theory treats the multiple scattering problem by considering the propagation of wave-energy in directed rays (essentially the short wavelength limit of wave optics) traversing a medium according to statistical considerations of the underlying turbid media. The MC technique employed in our simulation ‘samples’ this interpretation by propagating a finite number of quasi-plane-wave rays of wave energy²¹ which sum together to approach the solution to the underlying transport equation. Our problem thus reduces to defining the manner in which these rays propagate, and their associated phase perturbations due to the acousto-optic interaction.

Over $N-1$ sequential scattering events at $\vec{r}_j$, $\vec{r}_j$, $j = 1, \dots, N-1$, the phase increment of Eq. (6) becomes a summation $\sum_{j=1}^{N-1} \tilde{\phi}_d(\vec{r}_{j-1}, \vec{r}_j, t)$. With $\tilde{\xi}_s(\vec{r}_0, t) = 0$, we can write the phase increment upon incidence with a scatterer located at $\vec{r}$,

$$\tilde{\phi}_d(\vec{r}, t) = -n_0(\vec{k}_{0,j+1} - \vec{k}_{0,j}) \cdot \tilde{\xi}_s(\vec{r}, t), \quad (8)$$

where $\vec{k}_{0,j}$ and $\vec{k}_{0,j+1}$ represent the optical wavevectors prior to and following incidence with the scatterer, respectively. Thus, over a scattering path s involving N free paths and $N-1$ scattering events we may write the complex phase perturbation,

$$\tilde{\phi}_s(t) = \sum_{j=1}^{N-1} \tilde{\phi}_d(\vec{r}_j, t) + \sum_{j=1}^N \tilde{\phi}_n(\vec{r}_{j-1}, \vec{r}_j, t), \quad (9)$$

where j is a particular scattering event, $\vec{r}_0$ is the input location, and $\vec{r}_N$ is the location where the ray is detected. Taking the magnitude $m_s = |\tilde{\phi}_s|$ and phase angle $\varphi_s = \angle \tilde{\phi}_s(0)$ of this perturbation we may express the time varying phase increment as¹⁰

$$\phi_s(t) = \text{Re} \{ m_s \exp[j(\varphi_s - \omega_a t)] \}. \quad (10)$$

Each of our quasi-plane-wave rays will thus take the form $\tilde{E}_p(\vec{r}, t) = a_p(\vec{r}) \exp[j(\phi_{m,p}(\vec{r}, t) + \phi_{s,p}(t) - \omega_o t)]$ where a_p is some amplitude, $\phi_{m,p}$ some random phase due to the multiple scattering process for this ray, $\phi_{s,p}$ the AO phase shift for this ray, and ω_o the optical frequency.

2.2 Observables and Measures

Following multiple scattering in a turbid medium, the field variable $E(\vec{r}, t)$ will be a random variable fluctuating over time as effects such as Brownian motion alter the contributions from its constituent scattering paths. Nonetheless, we may investigate the effects of the AO modulation by means of the temporal-field autocorrelation function, $G_1(\tau)$, which is derived from the observable temporal-intensity autocorrelation function, $G_2(\tau)$, by the Siegert relationship,

$$|G_1(\tau)| = [G_2(\tau) - 1]^{1/2} = \left| \int_{-\infty}^{\infty} \tilde{E}(\vec{r}, t) \tilde{E}^*(\vec{r}, t - \tau) dt \right|. \quad (11)$$

In the context of our simulated quasi-plane-wave rays (herein photon packets),

$$G_1(\tau) \approx \sum_p \int_{-\infty}^{\infty} \tilde{E}_p(\vec{r}, t) \tilde{E}_p^*(\vec{r}, t - \tau) dt. \quad (12)$$

Neglecting Brownian motion (e.g., removing the time dependence of $\phi_m(\vec{r}, t)$) and under the assumption that the system is ergodic and wide-sense stationary, only those terms in the electric field which vary deterministically over time are retained in the expression for $G_1(\tau)$. Thus the integration can be performed by expansion of the exponentiated cosine via the Jacobi-Angers identity,²²

$$G_1(\tau) = \sum_p a_p^2 \left[J_0^2(m_{s,p}) + 2 \sum_{n=1}^{\infty} J_n^2(m_{s,p}) \cos(n\omega_a \tau) \right], \quad (13)$$

where J_n is the n 'th Bessel function of the first kind. We recognize that the modulation index, m_s , is proportional to the acoustic pressure amplitude [as demonstrated in Eqs. (2), (5), and (8)] and that the phase modulation expressed in Eq. (10) has caused the monochromatic optical input to be partially shifted in frequency. By the Wiener-Khinchin theorem the power spectrum of the detected field is given by the Fourier transform of the field autocorrelation function; this can be seen by inspection to consist of a DC component and numerous side-bands at multiples of the acoustic frequency, each scaled by a squared Bessel function of increasing order.

2.3 An Adjoint Method for AO

MC models of light transport are unable to directly simulate PSPD regimes which commonly occur in studies involving the use of coherent light, since the random walk process is unlikely ever to impinge directly upon the detection point. One approach to estimating the value of a quantity at a point with the MC method is the use of the adjoint method, as applied in the case of standard photon and neutron transport.^{23,24} We now expand this concept to the case of AO where energy will be distributed not only at the input optical frequency, but at sidebands shifted by multiples of the acoustic frequency.

A pencil beam incident normal to the boundary of the simulation domain at some point $\vec{r}_s$ illuminates the system with a given irradiance, $I(\vec{r}_s)$. We denote the probability of optical power at some point $\vec{r}_a$ in the k 'th sideband reaching some point $\vec{r}_b$ in the n 'th side band as $P_k^n(\vec{r}_b, \vec{r}_a)$. The probability

distribution of optical power at each sideband throughout the domain resulting from the source at $\vec{r}_s$ is given by the unitless normalized fluence rate $P_0^n(\vec{r}, \vec{r}_s) = \Phi^n(\vec{r})/I_0(\vec{r}_s)$, where $\Phi^n(\vec{r})$ is the fluence rate at a point $\vec{r}$ in the medium in the n 'th side-band. In replacing our detector with an adjoint source of equivalent directivity at $\vec{r}_d$, we may derive another set of optical power probability distributions $P_0^n(\vec{r}, \vec{r}_d)$.

The Chapman-Kolmogorov equation states²⁵

$$P(\vec{r}_d, \vec{r}_s) = \int P(\vec{r}_d, \vec{r})P(\vec{r}, \vec{r}_s)d^3\vec{r}. \quad (14)$$

In the case of transitions to and from different side bands we must sum over the relevant permutations:

$$P_0^n(\vec{r}_d, \vec{r}_s) = \int \sum_{k=0}^{\infty} P_k^n(\vec{r}_d, \vec{r})P_0^k(\vec{r}, \vec{r}_s)d^3\vec{r}. \quad (15)$$

We assume that contributions to the $(n+1)$ 'th side-band are negligible and thus express the detected exitance at the zeroth side-band,

$$W_0(\vec{r}_d) = I_0(\vec{r}_s) \int P_0^0(\vec{r}_d, \vec{r})P_0^0(\vec{r}, \vec{r}_s)d^3\vec{r}, \quad (16)$$

where by optical reciprocity $P_0^0(\vec{r}_d, \vec{r}) = P_0^0(\vec{r}, \vec{r}_d)$. The detected exitance at the first sideband is given by,

$$W_1(\vec{r}_d) = I_0(\vec{r}_s) \int P_0^1(\vec{r}_d, \vec{r})P_0^0(\vec{r}, \vec{r}_s) + P_1^1(\vec{r}_d, \vec{r})P_0^1(\vec{r}, \vec{r}_s)d^3\vec{r}. \quad (17)$$

We find $P_0^1(\vec{r}_d, \vec{r})$ by optical reciprocity, and make the approximation $P_1^1(\vec{r}_d, \vec{r}) \approx P_0^0(\vec{r}_d, \vec{r})$. Thus, the first side-band modulation depth [MD⁽¹⁾] detected at a point on the boundary is given by

$$\text{MD}^{(1)} \approx \frac{\int P_0^1(\vec{r}_d, \vec{r})P_0^0(\vec{r}, \vec{r}_s) + P_0^0(\vec{r}_d, \vec{r})P_0^1(\vec{r}, \vec{r}_s)d^3\vec{r}}{\int P_0^0(\vec{r}_d, \vec{r})P_0^0(\vec{r}, \vec{r}_s)d^3\vec{r}}. \quad (18)$$

3 Implementation

3.1 GPU Parallel Programming in the CUDA Framework

The design of an efficient GPU-accelerated parallel algorithm requires careful consideration of the underlying computational architecture. An extensive discussion of best practices is provided by the authors of the CUDA architecture.²⁶ Our implementation follows three broad strategies.

1. Exposing the parallelism of the algorithm: the problem at hand is inherently parallel as each photon packet is propagated independently of every other. In our algorithm, cases will occur when individual photon packets require different treatments. For example some may be reflected or refracted at an optical boundary, whereas others may be incident upon a scattering object. It is important to minimize this divergence since each execution branch will be serialized by the hardware.

2. Use of special-purpose memory: While such devices are capable of a very high arithmetic instruction throughput, this comes at the expense of the caching of memory. Various special-purpose memory spaces exist on the device which are optimized to store particular types of data.
3. Use of intrinsic arithmetic functions: CUDA exposes a number of functions which map directly to the hardware level; such instructions are considerably faster than the standard C-library implementations, but often have reduced accuracy. Judicious use of these functions (particularly the trigonometrical replacements) can considerably increase performance.

3.2 Architecture

The simulation program is built around the light-propagation and phase-accrual algorithms of the propagation kernel. Many thousands of instances of the kernel each execute in parallel on the GPU, the configuration and execution of which is managed by the driver. The driver consists of a MATLAB frontend to the simulation executable. The user configures the simulation by building a set of MATLAB variables and structures with the data required to define the simulation domain; this includes the geometry, medium properties, acoustic-field distribution and various simulation constants. The simulation is executed by passing this configuration structure to the MATLAB driver function. The driver validates the input data, constructs any dependent data, and calls the simulation executable. The executable uploads the data to the graphics card and repeatedly calls the kernel to propagate a certain number of photon packets until the desired number have been propagated. The output data is collected and stored in a MATLAB data file. The architecture of the simulation program is depicted in Fig. 1.

3.2.1 Geometry

Many MC simulation programs which support heterogeneous optical properties, including GPU-accelerated programs¹⁷ and those which consider AO,¹² implement heterogeneity by discretizing the domain into a voxelized grid. As a photon packet takes a step through multiple volume elements (voxels), each must be examined to ascertain if there has been a change in optical properties. While this approach has the benefit of being conceptually straightforward, representations of complex curved features require a very fine voxelization. Even with a fine voxelization, many of the problems discussed in Ref. 19 cannot be overcome.

A different approach is taken here whereby the simulation domain is represented by a tetrahedral mesh. This representation is typically encountered in applications of, e.g., the finite element method, but has also been applied to a CPU-based MC photon propagation problem.²⁷ By employing this method, photons move around within tetrahedral volume elements which subdivide the simulation geometry. On each photon step, we must check for intersection with any of the four facets of the current element; whilst we must still interrogate a dataset to determine the location of these facets, the amount of data to be retrieved, and the nature of the processing to be performed, is now fixed across each thread, thus avoiding issues of divergence and potentially reducing the amount of data which must be

Fig. 1 Architecture of the simulation program illustrating the parallel execution of many thousands of individual propagation kernels, each responsible for a single photon packet.

retrieved from memory relative to a voxelized solution (dependent upon the complexity of the model). Many of the advantages of this approach are discussed in Ref. 28.

The meshed geometry is represented by four sets of data. The first is the node list, which records the three-dimensional (3-D) Cartesian co-ordinate of every node, or vertex, in the mesh. The node list is used by the element connectivity matrix, which defines individual tetrahedral elements by indexing four such nodes. In addition to this data, we also provide a list of element neighbors that for each element provides an index into the four neighboring elements. Finally, the properties table indicates the medium of which an element is composed by providing an index into the medium array.

Each of the geometry tables is stored on the device in texture memory, which, whilst residing in the slow global memory, is spatially cached. This will convey significant advantages at the start of simulations when all photon packets reside in the same element.

3.2.2 Geometry intersection

The intersection of a line segment with a triangle is a well understood geometric operation. The most basic form of this test involves testing for intersection of the line with the plane defined by the three points of the triangle; many fast algorithms are available in the literature, one of the more popular being that of Möller and Trumbore.²⁹ When such tests are implemented

with finite precision, errors in the intersection tests may occur; this can often be tolerated in graphics rendering, for which such ray-tracing algorithms are typically developed. In this application the consequences of such an error is that a photon packet will end an iteration of the algorithm at a position which does not correspond to its current element; all subsequent intersection tests will fail and the photon will propagate indefinitely, or be fully absorbed within an incorrect element. This problem is exacerbated by the use of single-precision floating-point arithmetic within the kernel (the most efficient data type for current GPU architectures) rather than the double precision typically employed in a CPU-based algorithm.

A more robust approach is based upon the triple scalar product. In this case the majority of the tests are based upon the sign, rather than on the absolute value of the arithmetic operations. This approach is mathematically equivalent to the use of Plücker coordinates (as described by Fang²⁷), but avoids the associated complexity of analysis and notation. The signed volume of a tetrahedron defined by four vertices $\vec{p}_0$ – $\vec{p}_3$ is calculated by³⁰

$$V(\vec{p}_0, \vec{p}_1, \vec{p}_2, \vec{p}_3) = \frac{1}{6}[(\vec{p}_1 - \vec{p}_0) \times (\vec{p}_2 - \vec{p}_0) \cdot (\vec{p}_3 - \vec{p}_0)], \quad (19)$$

where $\cdot$ is the dot-product and $\times$ the cross-product operator. To determine if the ray defined by the current photon direction intersects the current tetrahedral element, we begin by

retrieving the four vertices of the element from the geometry cache. We test for intersection with each of the four facets of the tetrahedron by iterating over the four choices of the three nodes. We ignore the result of a test against the facet, if any, over which the photon packet previously reflected or refracted.

For each of the four facets we calculate three signed volumes $V(\vec{r}_a, \vec{r}_b, \vec{p}_1, \vec{p}_2)$, $V(\vec{r}_a, \vec{r}_b, \vec{p}_2, \vec{p}_0)$, and $V(\vec{r}_a, \vec{r}_b, \vec{p}_0, \vec{p}_1)$. In the case that the sign of each of these volumes is equal, the ray intersects the tetrahedron. To determine if the photon packet will intersect the tetrahedron given its current step size, we calculate the two signed volumes $V_a = V(\vec{p}_0, \vec{p}_1, \vec{p}_2, \vec{r}_a)$ and $V_b = V(\vec{p}_0, \vec{p}_1, \vec{p}_2, \vec{r}_b)$; in the case that the signs of these volumes are opposite, and that $V_a > 0$, intersection occurs at the parametric distance $V_a/(V_a + V_b)$. If any interactions occur, the smallest intersection distance is used to scale the current photon step size.

3.2.3 Acoustic field

As discussed in the theory, simulation of the AO modulation process requires knowledge of the local acoustic pressure amplitude, P_0 , the local acoustic wave-vector, $\hat{k}_a$, and the local acoustic phase offset relative to some point, ϕ_a . We pack this data into a four-element array by storing, for each point, $[P_0\hat{k}_{ax}, P_0\hat{k}_{ay}, P_0\hat{k}_{az}, \phi_a]$, where $\hat{k}_{ax}$ is the x component of the normalized wavevector. The magnitude of k_a is stored in constant memory to permit recovery of the wave-vector when required.

We store the acoustic field data in a 3-D texture to exploit hardware caching, hardware-based trilinear interpolation (interpolation of the acoustic-field dataset is only possible when the phase angle of the field is unwrapped; if discontinuities around 2π exist, interpolation will produce undesired effects), and normalized addressing. It is thus possible to read all needed parameters of the acoustic field at a given location in one texture access.

3.2.4 Medium

Each medium is represented by a structure stored in constant memory containing the absorption coefficient, μ_a , the scattering coefficient, μ_s , the anisotropy factor, g , the index of refraction, n , the piezo-optic coefficient, η , and the speed of sound, v_a .

3.2.5 Random number generator

On numerous occasions within the light-transport algorithm a uniformly distributed random number must be generated. This application uses an implementation by Alerstam¹⁶ of Marsaglia's multiply-with-carry random number generator.³¹ As described in the references, the algorithm requires only eight bytes of state information and has a period of greater than 2^{60} .

3.3 The Monte-Carlo Procedure

Many of the fundamental calculations performed are of a standard form for MC light-transport codes; Wang's description³² of his *de facto* standard code, MCML, will serve as a reference for these processes. We implement the following general procedure for light transport:

1. A photon packet is launched into the simulation domain according to the location and angular distribution of the source. In the case of mismatched refractive

indices, the photon packet-weight is reduced according to specular reflection [Eqs. (3.2)–(3.6) in Ref. 32].

2. The potential photon packet step size is sampled from the appropriate distribution [Eq. (3.13) in Ref. 32].
3. If the sampled step size will cause the the photon packet to intersect the boundary between two optically dissimilar regions (see Sec. 3.2.2), the step size is truncated to the distance to the boundary.
4. If the photon packet intersected an optical boundary:
 - a. The angle of incidence between the photon packet and boundary is calculated. With knowledge of the differing refractive indices, the angle of transmission is calculated.
 - b. The internal reflectance is calculated [Eq. (3.30) in Ref. 32] or set to unity in the case of total internal reflection. Reflection or transmission is determined [Eq. (3.31) in Ref. 32].
 - c. The next direction of the photon packet is calculated according to Snell's law, depending if reflection or transmission (refraction) has occurred.

Otherwise, the photon packet has been incident upon a scatterer:

- a. The next direction of the photon packet is calculated according to the Henyey-Greenstein phase function.³³
 - b. Some proportion of the current photon-packet weight is absorbed into the medium [Eq. (3.18) in Ref. 32]. The fluence at the absorption event is recorded on the nodes of the current tetrahedral element according to linear weightings of the distance of the photon packet to each node, with a proportion $J_0^2(m_s)$ distributed to the zeroth side band, and a proportion $2J_1^2(m_s)$ to the first. Following simulation the fluence rates are normalized by the total input optical power, and the voronoi volume of each node in the mesh.
 - c. The photon packet is subjected to the roulette procedure for variance reduction [Sec. (3.9) in Ref. 32].
5. The photon packet is moved according to the (potentially truncated) step size [Eq. (3.17) in Ref. 32], and the phase increment due to the AO modulation of the refractive index of the medium is calculated. The simulation offers two methods to calculate the index of refraction phase increment based upon Eq. (2).

In the first formulation we approximate Eq. (2) with a numerical integration over the path of the current photon-packet step. We divide the photon-packet path length into N subintervals of equal widths $h = |\vec{r}_b - \vec{r}_a|/N$ and employ the composite midpoint rule,

$$\tilde{\phi}_n(\vec{r}_a, \vec{r}_b, t) \approx k_0 n_0 \eta h \sum_{k=0}^{N-1} P_0(\vec{r}_k) \times \exp[j(\phi_a(\vec{r}_k) - \omega_a t)], \quad (20)$$

where $\vec{r}_k = \vec{r}_a + (k + 0.5)(\vec{r}_b - \vec{r}_a)/n$. The number of subintervals is defined at run-time as a particular number of points per acoustic wavelength and thus differs for each photon packet under propagation; the divergence of this approach leads to sub optimal performance. Furthermore, a large number of reads from the acoustic dataset are required to minimize the error in the numerical integral. This approach is still viable owing to the performance benefits of the parallel implementation and provides excellent accuracy irrespective of the nature of the underlying acoustic field.

An alternative approach to the calculation of the index of refraction phase increments is to assume that over the length of a single photon-packet step the acoustic field can be represented by a plane wave of given pressure amplitude, propagation direction and phase offset as defined at a point $\vec{r}_c$,

$$p_c(\vec{r}, t) = P_0(\vec{r}_c) \exp[j(\phi_a(\vec{r}_c) + \vec{k}_a(\vec{r}_c) \cdot (\vec{r} - \vec{r}_c) - \omega_a t)]. \quad (21)$$

This expression is substituted into Eq. (1) and the integration completed over the straight line path of the photon-packet step from $\vec{r}_a$ to $\vec{r}_b$,

$$\begin{aligned} \tilde{\phi}_n(\vec{r}_a, \vec{r}_b, t) &\approx \frac{k_0 n_0 \eta P_0(\vec{r}_c) \exp[j(\phi_a(\vec{r}_c) - \vec{k}_a(\vec{r}_c) \cdot \vec{r}_c - \omega_a t)]}{j \vec{k}_a(\vec{r}_c) \cdot (\vec{r}_b - \vec{r}_a)} \\ &\times (\exp[j(\vec{k}_a(\vec{r}_c) \cdot \vec{r}_b)] - \exp[j(\vec{k}_a(\vec{r}_c) \cdot \vec{r}_a)]). \end{aligned} \quad (22)$$

In the case that the acoustic wave-vector is a slowly changing function of space, this approximation can offer significant performance improvements over the numerical integration when the photon-packet step size is large compared to an acoustic wavelength. This is achieved due to reduced memory access, and its non divergent execution.

6. The photon-packet direction is updated according to the calculated next direction. If the photon packet is incident upon a scatterer, the phase increment due to the AO modulation of the scatterer location is calculated by application of Eq. (8).
7. If the photon packet has been completely absorbed, its propagation is terminated. If the photon packet has escaped the simulation domain, the location, weight, absolute path length and phase-modulation terms [m_s , ϕ_a of Eq. (10)] are recorded in a buffer residing in the device memory of the GPU, subsequently being transferred to the host and saved to disk. In all other cases the photon packet continues to propagate by returning to step two.

4 Results

4.1 Validation

The optical model of the simulation code has been validated against numerous analytical formulations based upon the DA to the radiative transfer equation (RTE). Simulations have been performed to compare the code to the total diffuse transmittance and total reflectance of van de Hulst,³⁴ the total reflectance as calculated in the multi layer diffusion model of Liemert and Kienle,³⁵ the temporal point-spread functions (TPSFs) presented by Martelli et al.,³⁶ and the internal fluence for an optically inhomogeneous sphere inside a semi-infinite slab as derived by Boas et al.³⁷ In each case, excellent agreement is found in regions where the DA is valid.

The optical model has also been compared against the output of MCML.³² A variety of simulations were performed in optically heterogeneous multi layered geometries. In all cases, excellent agreement was found in the total diffuse transmittance, total reflectance, and the radially resolved output.

Fewer analytical formulations are available with which to validate the AO model. Sakadžić and Wang⁸ present a set of figures [Figs. 2(a)–2(d)] demonstrating the MD⁽¹⁾ as derived by the analytical solution for AO in a slab geometry with plane-wave ultrasound. These figures have been recreated and demonstrate excellent agreement over each demonstrated permutation of the acoustic and optical set up. In subsequent work Sakadžić and Wang¹² employ their correlation diffusion equation to calculate the MD⁽¹⁾ within a highly scattering, optically homogeneous semi-infinite slab of width $x = 20$ mm, insonified by a cylinder of plane-wave ultrasound.

We recreate this simulation by generating a cuboid mesh of width $x = 20$ mm. The depth and height of the mesh ($y = 100$ mm, $z = 40$ mm) are chosen to minimize boundary effects. A cylinder of ultrasound is created of radius 3.175 mm. Ten million photon packets were launched into the domain at $x = 0$, $y = 10$ mm and allowed to propagate until they were completely absorbed into the medium or exited the domain via one of the boundaries. The simulation program recorded the internal fluence distribution on each node throughout the medium at the input optical frequency,

Fig. 2 IMD⁽¹⁾ [$\Phi^{(1)}(\vec{r})/\Phi^{(0)}(\vec{r})$] evaluated along the reflection ($x = 0$ mm, lower curve) and transmission ($x = 20$ mm, upper curve) surfaces of the slab in the plane of the source as calculated by the simulation program and the correlation diffusion equation. Excellent agreement is seen with the diffusion equation away from the source location in the reflection surface ($x = 0$, $y = 10$ mm).

Fig. 3 $\text{IMD}^{(1)} [\Phi^{(1)}(\vec{r})/\Phi^{(0)}(\vec{r})]$ for a cut through the slab in the plane of the source. The source is incident normal to the slab surface at $x = 0$ mm, $y = 10$ mm. The column of ultrasound is centered at $x = 10$ mm, $y = 0$ mm, as indicated by the white circle.

$\Phi^{(0)}(\vec{r})$, and at the first shifted side band, $\Phi^{(1)}(\vec{r})$. The internal first side-band modulation depth was calculated as $\text{IMD}^{(1)} = \Phi^{(1)}(\vec{r})/\Phi^{(0)}(\vec{r})$ prior to interpolation to a regular grid using standard MATLAB functions.

Figure 3 plots the internal first-side band-modulation depth $\text{IMD}^{(1)}$ inside the slab in the plane of the source. Figure 2 compares the MC simulation $\text{IMD}^{(1)}$ against that of the analytical solution along the transmission and reflection planes of the slab. Excellent agreement is found over the region where the DA is valid. Near to the source location we note a large dip in the $\text{IMD}^{(1)}$, as would be expected when the large unmodulated intensity is considered. The MC results agree excellently with those presented in the original work.¹²

In each of the simulations performed for the validation procedure, the explicit acoustic integration technique was employed since this is exact in the case of plane-wave insonification. Use of the numerical integration method with six points per wavelength produced identical results for this particular set of acoustic and optical properties.

4.2 Performance

To illustrate the performance of the simulation code, we compare its optical performance against a recent CPU mesh-based MC model, MMCM.²⁷ We execute the simulation of the inhomogeneous sphere in a scattering slab using the same mesh (MMCM Mesh 2) as used by Fang,²⁷ who reports a speed of 5.6 photon packets per millisecond (pp/ms) for a single thread, scaling to 21.2 pp/ms when executed over four threads.

Our GPU implementation processes 39.8 pp/ms while executing on a single nVidia GTX 250 graphics card.

The GPU implementation executes at almost eight times the speed of the single core CPU implementation, or almost twice that of the four-core parallel implementation of Fang.²⁷ The GPU implementation may be executed on a server incorporating multiple GPUs, where performance will scale linearly according to the number of devices available. While calculation of the AO phase shifts was disabled at run time, the overhead of variable definitions at compile time still has considerable performance implications owing to the number of threads which may execute simultaneously. Removal of the AO aspects of the simulation code should be expected to significantly improve performance if only an optical simulation is required.

Very little data has been published with which to compare the performance of the code in the AO case. We recently presented the results of a GPU-based slab geometry plane-wave AO code and compared these to a CPU implementation¹⁸ of the same. We now compare the performance of the current simulation code to that of both the CPU and GPU codes demonstrated previously.

A semi-infinite slab of depth 2 cm (represented in the current simulation code as a slab of lateral extent 10 cm) with optical properties $n_0 = 1.33$, $\mu_a = 0 \text{ cm}^{-1}$, $g = 0$, $\eta = 0.321$ and varying μ_s is illuminated on one surface and the AO $\text{MD}^{(1)}$ determined given insonification by a plane-wave acoustic field of displacement amplitude $A = 1.7 \text{ nm}$ and varying frequency. Table 1 details the speed of the simulations in numbers of photon packets per millisecond. The previously reported simulation figures are given in the column labeled GPU Orig., the speed of the current simulation is described for two frequencies and for both the numerical integration of phase of Eq. (20), and the explicit integration of Eq. (22).

To understand the resultant performance data, we must consider the limiting factors in each of the simulation algorithms. The performance of the original CPU and GPU codes is primarily a function of the scattering coefficient, μ_s , since this dictates the number of calculations which will be performed for a given photon packet's simulation (as discussed further in our earlier work¹⁸).

The new simulation code retains the same performance dependency while introducing two further factors which may affect performance.

1. If the average size of the mesh describing the simulation domain becomes smaller than the average step size of a photon packet (i.e., at low values of μ_s and with a dense mesh), photon-packet steps will routinely be truncated and geometry intersection tests will begin to dominate the processing workload. In this region of operation, the performance will be dependent upon the dimensions of the mesh.
2. When operating in the AO numerical integration mode a large photon-packet step size relative to the acoustic wavelength will result in multiple memory reads and sets of arithmetic operations to determine the value of the integral [i.e., $N > 1$ in Eq. (20)]. In this region, performance will depend upon the relative sizes of the average step size of a photon packet and the acoustic wavelength. This limitation does not affect the AO explicit integration mode, which currently performs only one calculation per photon packet step.

Table 1 Execution speed (photon packets/ms) of alternative acousto-optic simulation programs and configurations. Numbers in parentheses demonstrate the speed-up relative to the CPU implementation.

μ_s (cm ⁻¹)	Execution speed (photon packets/ms)					
	1 MHz				10 MHz	
	CPU ¹⁸	GPU Orig. ¹⁸	GPU (Num. Int.)	GPU (Explicit)	GPU (Num. Int)	GPU (Explicit)
5	9.43	1449.27 (×154)	232.15 (×25)	251.10 (×27)	71.85 (×8)	294.43 (×31)
10	4.60	1030.93 (×224)	166.17 (×36)	193.24 (×42)	62.08 (×13)	191.86 (×42)
20	2.25	568.18 (×253)	106.22 (×47)	109.80 (×49)	49.51 (×22)	109.24 (×49)
50	0.89	253.44 (×285)	44.54 (×50)	46.77 (×53)	30.38 (×34)	46.19 (×52)
100	0.44	87.57 (×199)	20.83 (×47)	20.64 (×47)	16.38 (×37)	19.52 (×44)
200	0.22	29.18 (×133)	8.71 (×40)	7.96 (×36)	8.29 (×38)	8.18 (×37)

The aforementioned qualities of the performance of the new simulation code are evident in the performance data of Table 1. Under insonification by ultrasound at a frequency of 1 MHz (when the acoustic wavelength is long compared to a mean free path) geometry intersection tests (and their associated memory access) dominate the execution time and the new simulation code performs similarly under both integration regimes.

The increased flexibility and associated complexity of the new simulation decrease the performance of the new code relative to the previously reported GPU implementation. The new code achieves a performance varying between 1/8th and 1/3th of the previously reported GPU code; as μ_s increases, the speeds of the two GPU codes approach one another, since fewer geometrical intersection tests are processed in the new code relative to the total computational workload.

Under 10 MHz insonification the explicit integration mode of the simulation maintains similar execution speed to that of 1 MHz insonification, though the numerical integration shows a significant relative degradation in performance at low values of μ_s .

In all but one of the examples investigated here the new simulation code is at least an order of magnitude faster than the CPU implementation; though as discussed, performance could be degraded by introducing a finer mesh, or in the case of the numerical integration mode, a higher frequency of insonification.

4.3 Sensitivity Mapping

In recent work, Gunadi and Leung³⁸ demonstrated the measurement of spatial sensitivity maps for the AO sensing technique. In this work an optical absorber of $\mu_a = 42$ cm⁻¹, $\mu_s = 12$ cm⁻¹, and dimensions $5 \times 5 \times 12$ mm was moved around a tank of intralipid of negligible absorption and equal scattering to that of the absorber. In the first case a transmission-mode geometry was investigated, with an optical source and detector separated by 35 mm. The absorber was moved around the central region in 1-mm steps over a 25×25 grid, in each case the detected optical intensity (in the absence of ultrasound) and the AO MD⁽¹⁾ (in the presence of ultrasound) were recorded. A reference measurement was taken in the absence of an absorber such that the sensitivity can be represented as $J = (X_{\text{ref}} - X)/X_{\text{ref}} \times 100\%$,

where X may be optical intensity or AO MD⁽¹⁾, and X_{ref} is the corresponding reference measurement.

We recreate this experiment in simulation though reduce the absorber size to a $1 \times 1 \times 12$ mm cube. The linear acoustic field is generated by distributing monopoles over the surface of a spherical cap of equal dimension to the transducer used in practice, and summing their contributions in the field. The amplitude is scaled to match that recorded by Gunadi and Leung.³⁸ The mesh and absorptive inclusion is built on the fly using the iso2mesh toolbox. Following simulation, the adjoint method described in Sec. 2.3 is used to calculate the detected MD⁽¹⁾. Figure 4(a) and 4(b) depict the optical and AO spatial sensitivity recovered by 1250 independent MC simulations, each of two million photon packets.

While the smaller absorptive inclusion leads to differences in the absolute sensitivity demonstrated by measurement and simulation, the form of the sensitivity distribution retains the same qualities as described by the experimental observations of Gunadi and Leung.³⁸ As would be expected, the optical sensitivity remains concentrated around the source and detector locations. The AO sensitivity demonstrates localization to the acoustic field and, more significantly, a very low level of sensitivity close to the source and detector.

The second case studied in Ref. 38 is a reflection-mode geometry with source-detector spacing of 30 mm. The ultrasound field is located centrally at a depth of 20 mm from the input and detection surface. The same optical absorber was moved around a scanning region of 30×20 mm, located centrally and starting at a depth of 5.5 mm from the input surface (owing to the dimensions of the experimental absorber). This experiment was recreated in simulation; Fig. 4(c) and 4(d) depict the optical and AO spatial sensitivity recovered by 1800 independent MC simulations of two million photon packets.

Once again, the absolute sensitivities of the simulation differ slightly to that of the experimental work owing to the differing absorber dimensions. The optical sensitivity has the expected 'banana' shape between the source and detector, with sensitivity rapidly declining at distances greater than around 10 mm. The AO sensitivity is concentrated at a depth of around 17 mm from the input/output plane; the peak relative sensitivity is not within the ultrasound focus region, as the lack of light at these depths outweighs the higher acoustic pressures. Close to the source/detector plane a region of negative AO sensitivity is found.

Fig. 4 Spatial sensitivity maps (% change from reference value). Inward arrows depict the source and outward arrows the detector location. The dotted ellipse illustrates the focus of the ultrasound region. (a) Transmission-mode optical sensitivity, (b) transmission-mode AO sensitivity, (c) reflection-mode optical sensitivity, and (d) reflection-mode AO sensitivity.

Negative sensitivity suggests that an absorptive perturbation has the effect of increasing the measurement quantity ($MD^{(1)}$). This can occur when the perturbation reduces the amount of ‘unmodulated’ light reaching the measurement position by an amount proportionally greater than it reduces light which has been modulated into the first side band. The form of the simulated sensitivities shows excellent agreement with the experimental work of Gunadi and Leung.³⁸

5 Conclusions and Future Work

The implementation of a highly parallel AO MC simulation code has been demonstrated. The optical model has been extensively validated against various analytical and numerical solutions, and the AO model against a complex AO analytical example.

The performance of the simulation code has been compared against CPU and GPU implementations of pre-existing optical and AO simulation codes. While our previous GPU implementation of an AO simulation code provides better performance, this is at the cost of significantly reduced flexibility in the nature of the simulation domain. The flexibility and performance of the code have also been demonstrated in the generation of spatial sensitivity maps which support experimental work currently under way.

The primary limitation of the current model is that it is designed purely for monochromatic acoustic excitation. The implementation of a polychromatic acoustic-field distribution would permit the simulation of modulated continuous-wave inputs, or pulsed ultrasound, as employed to gain spatial selectivity in the axis of the ultrasound. This could potentially be implemented in the style of Sakadžić’s voxelized CPU simulation.¹³

Further analysis of the relative validity of the explicit and numerical integration methods may allow the program to determine automatically the appropriate technique based upon the form of the supplied acoustic field.

It would also be advantageous to develop a theoretical perturbation model for AO which follows from the adjoint formulation presented in this work. This would then require only two MC simulations to be executed in order to provide an approximation to the actual spatial sensitivity of the technique. The efficiency of the simulation tool described here would permit such a perturbation model to be validated against multiple explicit sensitivity maps such as those described earlier.

Acknowledgments

The authors would like to thank Erik Alerstam for his GPU-accelerated CUDAMCML code and Sava Sakadžić for his solution to the correlation diffusion equation for the example provided in the validation of this work. Our thanks is also extended to Ben Cox, Jeremy Hebden and Simon Arridge for their useful discussions and assistance in this work. This research is funded by EPSRC (Grant No. EP/G005036/1).

References

1. P. Lai and J. R. McLaughlan et al., “Real-time monitoring of high-intensity focused ultrasound lesion formation using acousto-optic sensing,” *Ultrasound Med. Biol.* **37**(2), 239–252 (2011).
2. W. Leutz and G. Maret, “Ultrasonic modulation of multiply scattered light,” *Phys. B. Condens. Matter* **204**(1–4), 14–19 (1995).
3. G. Maret and P. E. Wolf, “Multiple light scattering from disordered media: the effect of brownian motion of scatterers,” *Z. Phys. B Condens. Matter* **65**(4), 409–413 (1987).

4. D. Pine and D. Weitz et al., "Diffusing wave spectroscopy," *Phys. Rev. Lett.* **60**(12), 1134–1158 (1988).
5. M. Kempe and M. Larionov et al., "Acousto-optic tomography with multiply scattered light," *J. Opt. Soc. Am. A* **14**(5), 1151 (1997).
6. G. D. Mahan, "Ultrasonic tagging of light: theory," *Proc. Natl. Acad. Sci. USA* **95**(24), 14015–14019 (1998).
7. L. Wang, "Mechanisms of ultrasonic modulation of multiply scattered coherent light: an analytic model," *Phys. Rev. Lett.* **87**(4), 043903 (2001).
8. S. Sakadžić and L. Wang, "Ultrasonic modulation of multiply scattered coherent light: an analytical model for anisotropically scattering media," *Phys. Rev. E* **66**(2), 026603 (2002).
9. S. Sakadžić and L. Wang, "Modulation of multiply scattered coherent light by ultrasonic pulses: an analytical model," *Phys. Rev. E* **72**(3), 036620 (2005).
10. F. J. Blonigen and A. Nieva et al., "Computations of the acoustically induced phase shifts of optical paths in acoustophotonic imaging with photorefractive-based detection," *Appl. Opt.* **44**(18), 3735–3746 (2005).
11. S. Sakadžić and L. V. Wang, "Correlation transfer equation for ultrasound-modulated multiply scattered light," *Phys. Rev.* **74**(3:2), 036618 (2006).
12. S. Sakadžić and L. Wang, "Correlation transfer and diffusion of ultrasound-modulated multiply scattered light," *Phys. Rev. Lett.* **96**(16), 163902 (2006).
13. S. Sakadžić and L. V. Wang, "Correlation transfer equation for multiply scattered light modulated by an ultrasonic pulse," *J. Opt. Soc. Am. A* **24**(9), 2797–2806 (2007).
14. A. Bratchenia and R. Molenaar et al., "Acousto-optic-assisted diffuse optical tomography," *Opt. Lett.* **36**(9), 1539–1541 (2011).
15. L. V. Wang, "Mechanisms of ultrasonic Modulation of multiply scattered coherent light: a monte Carlo model," *Opt. Lett.* **26**(15), 1191–1193 (2001).
16. E. Alerstam, T. Svensson, and S. Andersson-Engels, "Parallel computing with graphics processing units for high-speed Monte-Carlo simulation of photon migration," *J. Biomed. Opt.* **13**(6), 060504 (2008).
17. Q. Fang and D. A. Boas, "Monte-carlo simulation of photon migration in 3D turbid media accelerated by graphics processing units," *Opt. Express* **17**(22), 20178–20190 (2009).
18. T. S. Leung and S. Powell, "Fast Monte-Carlo simulations of ultrasound-modulated light using a graphics processing unit," *J. Biomed. Opt.* **15**(5), 055007 (2010).
19. T. Binzoni and T. S. Leung et al., "Light transport in tissue by 3D Monte Carlo: influence of boundary voxelization," *Comput. Meth. Programs Biomed.* **89**(1), 14–23 (2008).
20. V. Raman and K. S. Venkataraman, "Determination of the adiabatic piezo-optic coefficient of liquids," *Proc. Roy. Soc. A Mat.* **171**(945), 137–147 (1939).
21. A. J. Welch and M. J. Gemert, Eds., *Optical-Thermal Response of Laser-Irradiated Tissue*, Springer, Dordrecht, Netherlands (2011).
22. M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions: With Formulas, Graphs, and Mathematical Tables*, Applied Mathematics Series, Vol. **59**, Dover Publications, Dover, New York (1964).
23. I. Lux and L. Koblinger, *Monte Carlo Particle Transport Methods: Neutron and Photon Calculations*, CRC Press, Florida, USA (1991).
24. E. E. Lewis and W. Miller, *Computational Methods of Neutron Transport*, Wiley-Interscience, New York, USA (1993).
25. J. C. Schotland, J. C. Haselgrove, and J. S. Leigh, "Photon hitting density," *Appl. Opt.* **32**(4), 448–453 (1993).
26. NVidia, "CUDA C best practices guide," Technical Report March, nVidia (2011).
27. Q. Fang, "Mesh-based Monte Carlo method using fast ray-tracing in Plucker coordinates," *Biomed. Opt. Express* **1**(1), 165–175 (2010).
28. E. Margallo-Balbás and P. J. French, "Shape based Monte Carlo code for light transport in complex heterogeneous tissues," *Opt. Express* **15**(21), 14086–14098 (2007).
29. T. Möller and B. Trumbore, "Fast, minimum storage ray-triangle intersection," *J. Graph. GPU Game Tools* **2**(1), 21–28 (1997).
30. A. Kensler and P. Shirley, "Optimizing ray-triangle intersection via automated search," in *IEEE Symposium on Interactive Ray Tracing*, pp. 33–38, IEEE (2006).
31. G. Marsaglia, "Random number generators," *J. Mod. Appl. Stat. Meth.* **2**(1), 2–13 (2003).
32. L. Wang, "MCML Monte Carlo modeling of light transport in multi-layered tissues," *Comput. Meth. Program Biomed.* **47**(2), 131–146 (1995).
33. L. C. Henyey and J. L. Greenstein, "Diffuse radiation in the galaxy," *Astrophys. J.* **93**, 70 (1941).
34. H. C. van de Hulst, "Multiple light scattering," in *Tables, Formulas, and Applications*, Vol. **2** Academic Press, New York (1980).
35. A. Liemert and A. Kienle, "Light diffusion in a turbid cylinder II: layered case," *Opt. Express* **18**(9), 9266–9283 (2010).
36. F. Martelli and D. Contini et al., "Photon migration through a turbid slab described by a model based on diffusion approximation II. Comparison. ii. with Monte Carlo results," *Appl. Opt.* **36**(19), 4600–4612 (1997).
37. D. A. Boas and M. O'Leary et al., "Scattering of diffuse photon density waves by spherical inhomogeneities within turbid media: analytic solution and applications," *Proc. Natl. Acad. Sci. USA* **91**(11), 4887–4891 (1994).
38. S. Gunadi and T. S. Leung, "Spatial sensitivity of acousto-optic and optical (NIRS) sensing measurements," *J. Biomed. Opt.* **16**(12), 127005 (2011).