



Short note

A pseudospectral method for solution of the radiative transport equation

Samuel Powell^{a,*}, Ben T. Cox^b, Simon R. Arridge^c^a School of Biomedical Engineering & Imaging Sciences, King's College London, Becket House, London, SE1 7EU, UK^b Department of Medical Physics & Biomedical Engineering, University College London, London, WC1E 6BT, UK^c Centre for Medical Image Computing, University College London, London, WC1E 6BT, UK

ARTICLE INFO

Article history:

Received 28 April 2018

Received in revised form 6 October 2018

Accepted 16 January 2019

Available online 6 February 2019

Keywords:

Radiative transport

Light transport

Diffuse optics

Biomedical optics

Pseudospectral methods

ABSTRACT

The radiative transport equation accurately describes light transport in participating media, though analytic solutions are known only for simple geometries. We present a pseudospectral technique to efficiently compute numerical solutions to large-scale time-dependent and steady-state problems, with anisotropic scattering. A perfectly matched layer is proposed which allows this method to be applied in arbitrarily complex heterogeneous media. Our GPU-accelerated implementation of the technique is validated by comparison with a Monte-Carlo simulation, demonstrating excellent agreement.

© 2019 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The radiative transport equation (RTE) is a mono-energetic form of the Boltzmann transport equation, which models the propagation, attenuation and scattering of particles within participating media [1], and has been applied in numerous fields, including neutron transport, and atmospheric physics. In this work we consider the application of the RTE in the field of biomedical optics [2,3], where steady-state and time-domain solutions are sought to verify experiments and as part of model-based inverse problems.

The RTE, which can be derived directly from Maxwell's equations [4], describes the distribution of specific intensity, or radiance within a domain Ω ,

$$\left(\frac{1}{c} \frac{\partial}{\partial t} + \hat{\mathbf{s}} \cdot \nabla + \mu_t(\mathbf{r})\right) \phi(\mathbf{r}, \hat{\mathbf{s}}, t) = \mu_s(\mathbf{r}) \int_{S^{N-1}} p(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}') \phi(\mathbf{r}, \hat{\mathbf{s}}', t) d\hat{\mathbf{s}}' + q(\mathbf{r}, \hat{\mathbf{s}}, t) \quad (1)$$

where the change in radiance $\phi(\mathbf{r}, \hat{\mathbf{s}}, t)$ at a point $\mathbf{r} \in \Omega$ in direction $\hat{\mathbf{s}}$ at time t , is given as a balance of energy lost from an attenuation term $\mu_t = \mu_a + \mu_s$ accounting for absorption (μ_a) and out-scattering (μ_s), inwards scattering from \mathbf{s}' to \mathbf{s} , and any sources $q(\mathbf{r}, \hat{\mathbf{s}}, t)$. In the biomedical imaging modalities of interest here, the Henyey–Greenstein phase function [5] is typically employed to describe the anisotropic scattering process,

* Corresponding author.

E-mail address: s.powell@kcl.ac.uk (S. Powell).

$$p(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}') = \frac{1}{4\pi} \frac{1 - g^2}{(1 + g^2 - 2g(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}'))^{3/2}}. \quad (2)$$

Boundary conditions for the RTE specify that $\phi(\mathbf{m}, \hat{\mathbf{s}}, t) = 0$ for $\hat{\mathbf{s}} \cdot \hat{\mathbf{n}} < 0$, where $\hat{\mathbf{n}}$ is the outward normal to the boundary of the domain at $\mathbf{m} \in \partial\Omega$: that is to say that in the absence of sources, there is no component of the radiance inwards across the boundary.

The high dimensionality of the field variable $\phi(\mathbf{r}, \hat{\mathbf{s}}, t) \in \mathcal{R}^N \times \mathcal{S}^{N-1}$, and the presence of non-local spatial and angular operators, is such that analytical solutions to the RTE have been found only for homogeneous infinite and semi-infinite geometries [6], and layered media [7].

Numerical solutions to the RTE in arbitrary three-dimensional geometries are typically sought by stochastic Monte-Carlo (MC) estimates. This approach converges to the solution in the limit of an infinite number of trials, and is inherently parallelisable. Despite the availability of a number of high performance codes, it remains a challenge to generate high quality estimates at distances far from a source, where statistical noise dominates. Furthermore, practical implementations require the use of variance reduction techniques which produce complex geometry dependent noise statistics which prevents modelling of the noise statistics as part of, e.g., an image reconstruction procedure.

Deterministic methods can be categorised by how they treat the angular discretisation.

- Discrete ordinate (S_N) methods discretise the angular basis over a set of vectors equally spaced over the unit-circle or sphere. S_N methods suffer from ray effects and may require stabilisation which affects both the performance and accuracy of their solutions.
- Spherical harmonic discretisations (P_N methods) expand the solution in the natural orthogonal basis over the sphere, and truncate the series to allow numerical solution. P_N methods can suffer from wave-like distortion owing to the truncation of the angular basis, but perform well in highly scattering media such as biological tissue.

The use of the RTE within biomedical optics extends across a number of modalities. Low-resolution time-resolved modalities such as diffuse optical tomography (DOT) require time-dependent solutions on picosecond temporal scales in domains c. 100 mm on a side, with feature sizes of a few millimeters. High-resolution hybrid modalities such as photoacoustic tomography require steady-state solutions in smaller domains c. 10 mm on a side, but may require sub-millimeter resolution. In all cases, domains are scattering dominated, with scattering coefficients c. 10 mm^{-1} compared to absorption coefficients of c. 0.01 mm^{-1} , at the near-infrared wavelengths of interest. In this diffuse regime, P_N methods are particularly appropriate as the smooth radiance fields which result are accurately captured in approximations of a lower order than would be required in an S_N method.

Discretisation in space can be achieved with the finite-element method [8–10], though the large size of the resultant system matrices has limited the practical application of this method. Alternatively, finite-difference [11] and finite-volume [12] methods have been proposed, but the inherent derivative operations require fine discretisation and the use of higher order derivative stencils, leading to limitations in their application in large domains. Mesh-free techniques have also been proposed [13]. There exists a significant body of work on solutions to the transport equation in other fields, such as neutron transport [14]. Here S_N methods are more appropriate, and computational performance requirements can be relaxed since the models are not typically employed as part of iterative inverse solvers.

The lack of accurate and computationally efficient deterministic methods often leads to the application of the diffusion approximation (DA) to the RTE. The DA fails in regions of low scattering, and where the gradient of the energy density is large. Furthermore, the DA is parabolic (non-causal) in contrast with the hyperbolic RTE from which it is derived. Hybrid transport-diffusion methods have been proposed to balance the computational challenges of transport solutions with the approximations inherent in the DA [15,16].

Whilst these differences have until recently been tolerated in the context of diffuse optics, new developments in high-density, time-domain, and hybrid diffuse optical methods require the use of the RTE to achieve their full potential.

2. The pseudospectral method

Pseudo-spectral methods are typically applied to solve partial differential equations which contain point-wise potential operators. The technique exploits the $\mathcal{O}(N \log N)$ performance of the Fast Fourier Transform (FFT) to take spatial derivatives in the spectral domain (k-space) before applying an inverse transform to apply the potential operators in Cartesian space: thus, all operators are applied in the space in which they are diagonalised, and a far coarser spatial discretisation can be employed, limited only by the Nyquist criterion. Such techniques thus require far less memory and computational effort than their finite-difference or finite-element counterparts.

Pseudo-spectral schemes can be constructed for the RTE using both the S_N and P_N methods, however the elegance of our proposed technique is to exploit the fact that the scattering operator of the RTE is also diagonalised in the P_N approximation, and that directional derivative term forms a recurrence relation over the spherical harmonics which is trivial to compute. We thus operate solely in the angular-spectral domain.

To derive our method we expand the radiance and source terms in a Spherical Harmonic basis,

$$\phi(\mathbf{r}, \hat{\mathbf{s}}, t) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} \sqrt{\frac{2\ell+1}{4\pi}} \psi_{\ell}^m(\mathbf{r}, t) Y_{\ell}^m(\hat{\mathbf{s}}), \quad (3)$$

$$q(\mathbf{r}, \hat{\mathbf{s}}, t) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} \sqrt{\frac{2\ell+1}{4\pi}} q_{\ell}^m(\mathbf{r}, t) Y_{\ell}^m(\hat{\mathbf{s}}), \quad (4)$$

and employ the addition theorem for spherical harmonics to restate the phase function,

$$p(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}') = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} g^{\ell} \bar{Y}_{\ell}^m(\hat{\mathbf{s}}') Y_{\ell}^m(\hat{\mathbf{s}}). \quad (5)$$

Equations (3) through (5) are substituted into equation (1), and the result is projected into the spherical harmonic basis by taking the inner product with $\bar{Y}_{\ell}^m(\hat{\mathbf{s}})$. Following some algebra, this yields a set of coupled first order PDEs [17],

$$\begin{aligned} \left(\frac{1}{c} \frac{\partial}{\partial t} + \mu_t \right) \psi_{\ell}^m + \frac{1}{2\ell+1} \left(\frac{\partial}{\partial z} [\alpha_{\ell+1}^m \psi_{\ell+1}^m + \alpha_{\ell}^m \psi_{\ell-1}^m] - \frac{1}{2} \left(\frac{\partial}{\partial x} - i \frac{\partial}{\partial y} \right) [\beta_{\ell}^m \psi_{\ell-1}^{m-1} - \beta_{\ell+1}^{-m+1} \psi_{\ell+1}^{m-1}] \right. \\ \left. - \frac{1}{2} \left(\frac{\partial}{\partial x} + i \frac{\partial}{\partial y} \right) [-\beta_{\ell}^{-m} \psi_{\ell-1}^{m+1} + \beta_{\ell+1}^{m+1} \psi_{\ell+1}^{m+1}] \right) = \mu_s g^{\ell} \psi_{\ell}^m + q_{\ell}^m, \end{aligned} \quad (6)$$

where the terms $\alpha_{\ell}^m = \sqrt{(l+m)(l-m)}$ and $\beta_{\ell}^m = \sqrt{(l+m)(l+m-1)}$ are related to the Clebsch–Gordan coefficients [8,17].

To implement our pseudo-spectral scheme we perform an N dimensional spatial Fourier transform of the radiance distribution,

$$\mathcal{F}[\phi(\mathbf{r}, \hat{\mathbf{s}}, t)](\mathbf{k}, \hat{\mathbf{s}}, t) = u(\mathbf{k}, \hat{\mathbf{s}}, t) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} \sqrt{\frac{2\ell+1}{4\pi}} Y_{\ell}^m(\hat{\mathbf{s}}) \hat{\psi}_{\ell}^m(\mathbf{k}, t) \quad (7)$$

where

$$\hat{\psi}_{\ell}^m(\mathbf{k}, t) = \int \psi_{\ell}^m(\mathbf{r}, t) \exp(-i\mathbf{k} \cdot \mathbf{r}) \, d\mathbf{r}, \quad (8)$$

and $\mathbf{k} = \{k_x, k_y, k_z\}$ is the k -space wave-vector. Inserting this result into equation (6) results in a set of coupled ODEs,

$$\left(\frac{1}{c} \frac{\partial}{\partial t} + \mu_t \right) \psi_{\ell}^m + \mathcal{F}^{-1}[\hat{\psi}_{\ell}^m] = \mu_s g^{\ell} \psi_{\ell}^m + q_{\ell}^m, \quad (9)$$

where \mathcal{F}^{-1} denotes the inverse spatial Fourier transform, and

$$\begin{aligned} \hat{\psi}_{\ell}^m = \frac{1}{2\ell+1} \left(-ik_z [\alpha_{\ell+1}^m \hat{\psi}_{\ell+1}^m + \alpha_{\ell}^m \hat{\psi}_{\ell-1}^m] - \frac{1}{2} (-ik_x - k_y) [\beta_{\ell}^m \hat{\psi}_{\ell-1}^{m-1} - \beta_{\ell+1}^{-m+1} \hat{\psi}_{\ell+1}^{m-1}] \right. \\ \left. - \frac{1}{2} (-ik_x + k_y) [-\beta_{\ell}^{-m} \hat{\psi}_{\ell-1}^{m+1} + \beta_{\ell+1}^{m+1} \hat{\psi}_{\ell+1}^{m+1}] \right). \end{aligned} \quad (10)$$

2.1. Discretisation, truncation, transient and steady-state solution

To solve the RTE numerically we form a P_N approximation by closing the recurrence under the assumption that $\psi_{\ell}^m = 0$ for $l > N$, leading to $(N+1)^2$ coupled first-order ODEs. The radiance distribution is sampled in space and the continuous Fourier transform replaced with its discrete counterpart. Finally, we sample in the temporal domain and implement a numerical approximation to the temporal derivative. For brevity, we choose here to employ a first order explicit time-stepping scheme,

$$\psi_{\ell,t+1}^m = \psi_{\ell,t}^m + c\Delta t \left(\mathcal{F}^{-1}[\hat{\psi}_{\ell,t}^m] + (\mu_t - \mu_s g^{\ell}) \psi_{\ell,t}^m \right) + q_{\ell,t}^m. \quad (11)$$

The use of an explicit temporal discretisation provides conditional stability, requiring that a suitable choice of time-step Δt be employed. We have found that time-steps of the order of 100 fs provide stability in the explicit time-stepping scheme across a wide range of physiologically relevant parameters. Typical time-domain measurements acquired using time-correlated single photon counting (TCSPC) hardware provide measurements with time-bins of c. 800 fs. Thus, for time-domain measurements the use of an explicit scheme offers reasonable efficiency.

To compute a steady-state solution, this time-domain solution can be numerically integrated, though the small step-size required for stability of the explicit scheme makes this approach inefficient. Alternatively, the time-domain form can be

integrated analytically, whence the pseudo-spectral method provides an implicit operator representing the underlying transport equation. Krylov subspace iterative methods can then be employed to find a steady-state solution through repeated application of the implicit operator. This technique is related to the method of exponential integrators, a topic to which we shall return in the discussion.

2.2. Boundary conditions

The primary limitation of pseudo-spectral methods is that they imply periodic boundary conditions (the exact form of which depends upon the choice of transform). In many applications a-periodic boundary conditions are enforced by the development of a *perfectly matched layer* (PML), which attenuates outgoing radiation before it reaches the boundary of the *computational* domain. Since the transport operator in the RTE is advective (that is to say that the action of transport on functions in the domain is to translate only along the directions $\hat{\mathbf{s}}$), a PML can be implemented by the simple addition of an additional isotropic loss term outside of the real domain. Further, to adhere to the specified boundary conditions, which state that there is no radiation inwards across the boundary of the domain, we must set the scattering coefficient $\mu_s = 0$ in those regions which do not correspond the real domain interest. A full analysis of such boundary conditions was recently provided in [18]. We thus solve,

$$\begin{aligned} \left(\frac{1}{c} \frac{\partial}{\partial t} + \mu_t\right) \psi_\ell^m + \mathcal{F}^{-1}[\dot{\psi}_\ell^m] &= \mu_s g^\ell \psi_\ell^m + q_\ell^m & \mathbf{r} \in \Omega, \\ \left(\frac{1}{c} \frac{\partial}{\partial t} + \mu_b\right) \psi_\ell^m + \mathcal{F}^{-1}[\dot{\psi}_\ell^m] &= 0 & \mathbf{r} \in \Omega_c \setminus \Omega, \end{aligned} \quad (12)$$

where μ_b is the PML attenuation term, Ω_c is the full computational domain, and Ω is the real domain of interest. The exact spatial distribution of μ_b should be chosen to ensure that the solution is attenuated to a sufficient degree on the boundary of Ω_c . Choice of μ_b is straightforward since within Ω_c the radiance decays at a rate $\exp(-\mu_b |r_b|)$, where r_b is the distance from the boundary in a given direction. It is also desirable for numerical reasons to choose a distribution which leads to a smooth decay of the radiance, in order to limit the introduction of higher spatial frequencies in the spectral representation.

3. Results

The pseudospectral algorithm was implemented in the Julia programming language [19]. The application of the transport operator occurs in three stages.

1. A three-dimensional FFT is applied over the spatial dimensions of the complex valued input radiance field. The directional gradient operator ($\mathbf{s} \cdot \nabla$) is then applied to the resulting k-space coefficients according to eq. (10). The solution is returned from k-space through application of an inverse transform.
2. The potential operators ($\mu_t - \mu_s g^\ell$) are applied to the input radiance field, and the result from (1) is added. This is a simple pointwise operation, and concludes application of the transport operator.
3. If required, time stepping is applied according to eq. (11) which updates the input radiance field in-place.

The algorithm defaults to running on the host CPU, where the FFTW library is used to implement the FFT. To improve execution speed we have also implemented the algorithm as a set of custom CUDA kernels which execute on the GPU, in this configuration the cuFFT library is used to implement the FFT.

All simulations were executed on an nVidia Tesla K40 GPU driven by a 32-core Xeon workstation. In the numerical demonstrations which follow, we in each case compare the results of our pseudo-spectral method with a Radiance Monte-Carlo technique we developed previously [20,21].

3.1. Transient solution without PML

The basic algorithm can be validated without the use of the PML by evaluating the time-domain internal radiance distribution over a period before the radiance distribution has propagated to the boundary.

We consider a cube 50 mm on a side of background parameters $\mu_a = 0.01 \text{ mm}^{-1}$, $\mu_s = 1.0 \text{ mm}^{-1}$, and an anisotropy factor of $g = 0.9$. The region $x \leq 15 \text{ mm}$ has a higher scattering coefficient $\mu_s = 5.0 \text{ mm}^{-1}$. The domain is discretised into a $64 \times 64 \times 64$ grid of points, and the pseudo-spectral method is used to solve the system under a 21-degree spherical harmonic (P_{21}) approximation.

The sharp onset of the scattering perturbation, and otherwise low scattering, provide an artificially challenging numerical scenario for the pseudo-spectral solver, promoting the propagation of higher angular harmonics and generating sharp features in the resultant fluence field.

A time-step of 0.1 ps was chosen by inspection, and 400 steps were taken such that the initial radiance distribution, consisting of an Gaussian of isotropic angular profile, with $1/e$ diameter 1.5 mm, has time to propagate close the boundary of the domain. The simulation took 70 s to complete.

The same parameters were also employed in a Monte-Carlo solver of the RTE, in which 5×10^8 packets were propagated, with a kernel run-time of 269 s.

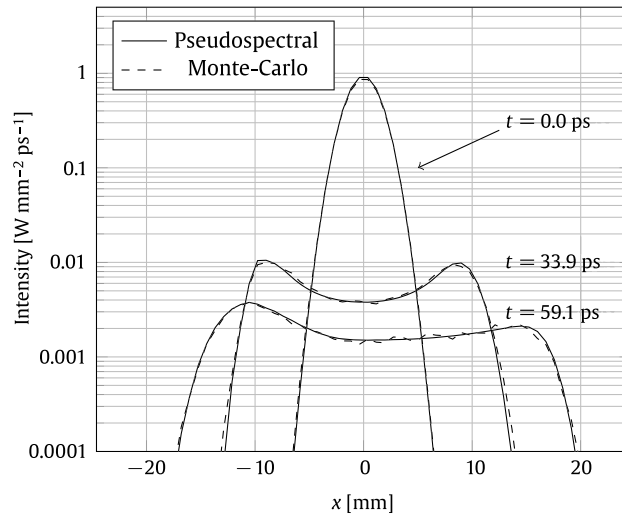


Fig. 1. Fluence along the line $y = z = 0$ mm at three points in time, pseudo-spectral method (solid) vs. Monte-Carlo (dots).




Fig. 2. Fluence distribution in the domain at $z = 25$ mm and $t = 65$ ps, pseudo-spectral method (left) vs. Monte-Carlo (right).

Fig. 1 demonstrates the evolution of the radiance field during propagation over a line through the domain. Fig. 2 shows the resultant fluence through the domain after 60 ps, when the simulation was terminated. The initial field spreads outwards through the region of low-scattering close to the speed of light, forming a circular disc in the plane. As the radiance encounters the region of increased scattering a characteristic region of back-scattering causes a peak at the interface, and the further transport is inhibited and the field becomes diffuse.

An exact comparison is challenging since the time-domain Monte-Carlo solution is inherently averaged spatially and temporally (whereas the pseudospectral method is a collocation technique, with an exact solution at points in space and time). Further reducing the spatial and temporal discretisation of the MC solution leads to an unacceptable reduction in SNR, highlighting the value of the proposed technique. Nonetheless, excellent agreement is found between the two solutions.

3.2. Steady-state internal fluence with PML

The effect of the PML can be evaluated by considering the steady-state fluence in the domain resulting from a given source distribution.

The real domain consists of a cuboid of 35.5 mm on a side with homogeneous optical parameters $\mu_a = 0.01$ mm⁻¹, $\mu_s = 10.0$ mm⁻¹, and anisotropy factor $g = 0.95$. The real domain is embedded into a computation domain 50 mm on a side with a PML attenuation term following a half-cosine profile rising from 0.01 mm⁻¹ at the end of the real domain to 1.0 mm⁻¹ at the boundary. The domain is discretised into a $64 \times 64 \times 64$ grid of points, and the pseudo-spectral method is used to solve the system under a 7-degree spherical harmonic (P_7) approximation.

A time step of 0.04 ps was chosen by inspection, and 25×10^3 time steps were integrated to ensure the domain had reached steady state conditions. The source term was an isotropic Gaussian distribution, with $1/e$ diameter 1.5 mm, offset from the centre of the domain by 8 mm in the x -axis. The simulation took 223 s to complete.

The same parameters were also employed in a Monte-Carlo solver of the RTE, in which only the real domain was simulated and a matched boundary condition employed. The MC simulation employed 1×10^8 packets, with a kernel run-time of 16.1 s.




Fig. 3. Steady-state fluence distribution in the domain across the x dimension at three points in y , at $z = 0$ mm, pseudo-spectral method (solid) vs. Monte-Carlo (dots).

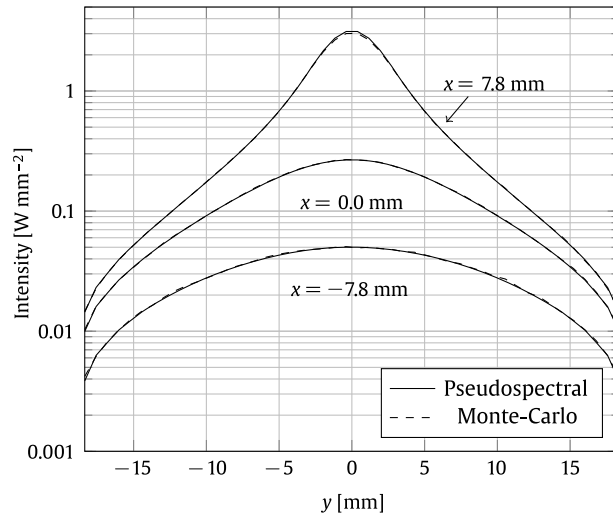


Fig. 4. Steady-state fluence distribution in the domain across the y dimension at three points in x , at $z = 0$ mm, pseudo-spectral method (solid) vs. Monte-Carlo (dots).

Figs. 3 and 4 demonstrate the resultant steady state fluence for both the pseudo-spectral and Monte-Carlo techniques. In each case the characteristic logarithmic decay of the fluence is evident away from the source and boundary, with an increased decay towards the boundary. The index-matched boundary conditions imply a non-zero fluence at the boundary the domain, evident in each simulation.

Excellent agreement is shown between the pseudo-spectral method and the Monte-Carlo solution.

4. Discussion & conclusions

We have introduced a new technique to solve the radiative transport equation in arbitrary heterogeneous media. The technique has been validated against a stochastic Monte-Carlo solver, demonstrating excellent agreement.

The use of a pseudo-spectral method allows a compact representation of the radiance field with sparser sampling than required in a finite-difference or finite-element scheme. The method is inherently parallelisable: whilst our implementation only uses a single GPU, the cuFFT library and CUDA kernels implement the algorithm can be readily extended to run across a number of devices with a minimum of inter-device communication. Our implementation is limited by the performance of the FFT, and thus the method demonstrates $\mathcal{O}(N \log N)$ performance scaling.

The pseudospectral technique is particularly suited to biomedical imaging where inherent scattering leads to (angularly) smooth solutions which can be simulated with modest angular discretisations (e.g., $N_l \leq 7$). As we have demonstrated, the

technique can readily support propagation in low-scattering regions when a higher degree approximation is employed. The absence of noise and deterministic performance is such that the technique could be employed as part of a conventional image reconstruction method in modalities such as diffuse optical tomography, quantitative photoacoustic tomography, and ultrasound modulated optical tomography.

The simplicity and flexibility of the technique suggests numerous avenues for further research. A particularly interesting avenue of investigation is the optimisation of the pseudo-spectral method for the computation of steady-state solutions. We noted earlier that the use of a Krylov solver to form such solutions directly from the integrated transport operator was an alternative to the relatively inefficient numerical integration of the time-domain solution. At present the computational performance of this approach is limited as Krylov methods are typically implemented on the CPU, necessitating significant overhead in memory transfers between the GPU and host. Overcoming this would yield an attractive technique. We have also successfully employed exponential integrators to implement arbitrary time steps in homogenous media, allowing for more efficient numerical integration.

Acknowledgements

S. Powell acknowledges support from Royal Academy of Engineering fellowship RF1516\15\33. S. Powell and S.R. Arridge acknowledge support from EPSRC grant EP/N032055/1. The authors declare no relevant conflicts of interests in this work.

References

- [1] K.M. Case, P.F. Zweifel, *Linear Transport Theory*, Addison-Wesley Series in Nuclear Engineering, Addison-Wesley, 1967.
- [2] S. Chandrasekhar, *Radiative Transfer*, Dover Publications, New York, NY, 1960.
- [3] A. Ishimaru, *Wave Propagation and Scattering in Random Media*, Wiley-IEEE Press, 1999.
- [4] M.I. Mishchenko, *J. Quant. Spectrosc. Radiat. Transf.* 101 (2006) 540.
- [5] L.G. Henyey, J.L. Greenstein, *Astrophys. J.* 93 (1941) 70.
- [6] A. Liemert, A. Kienle, *Sci. Rep.* 3 (2013).
- [7] A. Liemert, D. Reitzle, A. Kienle, *Sci. Rep.* 7 (2017) 3819.
- [8] S.R. Arridge, T. Tarvainen, M. Schweiger, P. Surya Mohan, A. Pulkkinen, *J. Comput. Phys.* 230 (2011) 7364.
- [9] T. Tarvainen, M. Vauhkonen, V. Kolehmainen, J. Kaipio, *Appl. Opt.* 44 (2005) 876.
- [10] M.A. Badri, P. Jolivet, B. Rousseau, Y. Favennec, *J. Comput. Phys.* 360 (2018) 74.
- [11] M. Sakami, K. Mitra, P.F. Hsu, *J. Quant. Spectrosc. Radiat. Transf.* 73 (2002) 169.
- [12] F. Aslanaj, S. Fumeron, *J. Biomed. Opt.* 17 (2012) 075007.
- [13] M. Kindelan, F. Bernal, P. González-Rodríguez, M. Moscoso, *J. Comput. Phys.* 229 (2010) 1897.
- [14] T.A. Wareing, J.M. McGhee, J.E. Morel, S.D. Pautz, *Nucl. Sci. Eng.* 138 (2001) 256.
- [15] J.D. Densmore, K.G. Thompson, T.J. Urbatsch, *J. Comput. Phys.* 231 (2012) 6924.
- [16] M. Roger, C. Caliot, N. Crouseilles, P.J. Coelho, *J. Comput. Phys.* 275 (2014) 346.
- [17] S.R. Arridge, *Inverse Probl.* 15 (1999) R41.
- [18] H. Egger, M. Schlottbom, arXiv:1802.08305, 2018.
- [19] J. Bezanson, A. Edelman, S. Karpinski, V.B. Shah, arXiv:1411.1607v4, 2014.
- [20] R. Hochuli, S. Powell, S.R. Arridge, B.T. Cox, *SPIE BiOS* 9323 (2015) 93231P.
- [21] S. Powell, R. Hochuli, S.R. Arridge, in: B.J. Tromberg, A.G. Yodh, E.M. Sevick-Muraca, R.R. Alfano (Eds.), *Optical Tomography and Spectroscopy of Tissue XII*, 2017, 100590W.