Monte Carlo fluence simulation for prospective evaluation of interstitial photodynamic therapy treatment plans

Jeffrey Cassidy[a], Vaughn Betz[a] and Lothar Lilge[b]

[a]Electrical and Computer Engineering, University of Toronto, Canada;
[b]Princess Margaret Cancer Centre, Toronto, Canada

ABSTRACT

Photodynamic therapy (PDT) delivers a localized cytotoxic dose that is a function of tissue oxygen availability, photosensitive drug concentration, and light fluence. Providing safe and effective PDT requires an understanding of all three elements and the physiological response to the radicals generated. Interstitial PDT (IPDT) for solid tumours poses particular challenges due to complex organ geometries and the associated limitations for diffusion theory based fluence rate prediction, in addition to restricted access for light delivery and dose monitoring.

As a first step towards enabling a complete prospective IPDT treatment-planning platform, we demonstrate use of our previously developed FullMonte tetrahedral Monte Carlo simulation engine for modelling of the interstitial fluence field due to intravesicular insertion of fibre light sources. The goal is to enable a complete treatment planning and monitoring workflow analogous to that used in ionizing radiation therapy, including plan evaluation through dose-volume histograms and algorithmic treatment plan optimization.

FullMonte is to our knowledge the fastest open-source tetrahedral MC light propagation software. Using custom hardware acceleration, we achieve 4x faster computing with 67x better power efficiency for limited-size meshes compared to the software. Ongoing work will improve the performance advantage to 16x with unlimited mesh size, enabling algorithmic plan optimization in reasonable time.

Using FullMonte, we demonstrate significant new plan-evaluation capabilities including fluence field visualization, generation of organ dose-volume histograms, and rendering of isofluence surfaces for a representative bladder cancer mesh from a real patient. We also discuss the advantages of MC simulations for dose-volume histogram generation and the need for online personalized fluence-rate monitoring.

Keywords: Monte Carlo, PDT, Photodynamic therapy, interstitial, dose-volume histogram

1. INTRODUCTION

In addition to clinical trials discussed below, a number of researchers[1,2] have articulated goals for interstitial photodynamic therapy (IPDT) treatment planning and dosimetry workflows. Superficial applications[3] can generally be delivered based on simple rules regarding surface irradiance and photosensitizer dose, whereas interstitial applications particularly must cope with optical inhomogeneity in three dimensions which can greatly impact the fluence rate inside the tissue. The situation may be further complicated by close proximity to organs at risk, and by the limited degree of access to the region for light delivery and monitoring. Like others, our concept[4] shown in Fig 1, which is based heavily on radiotherapy workflows, relies on numerical fluence simulation to mitigate some of the uncertainty in delivery. It proposes new processes to account for sources of variability in PDT treatment including source placement error and variability of the patient's personal optical properties, which can be quite large. Successfully incorporating all of these factors into a treatment planning and online-adjustment workflow will lead to truly personalized interstitial PDT delivery.

Previous efforts in IPDT have generally used the diffusion approximation to simplify light propagation calculations. For infinite homogeneous geometries, analytic solutions using the diffusion approximation[5] exist and have been used for prostate PDT dosimetry.[1] A more sophisticated approach uses the same physical approximation, but permits differing geometry through use of the finite-element method. Prostate cancer trials conducted

Send correspondence to JC via E-mail: jeffrey.cassidy@mail.utoronto.ca

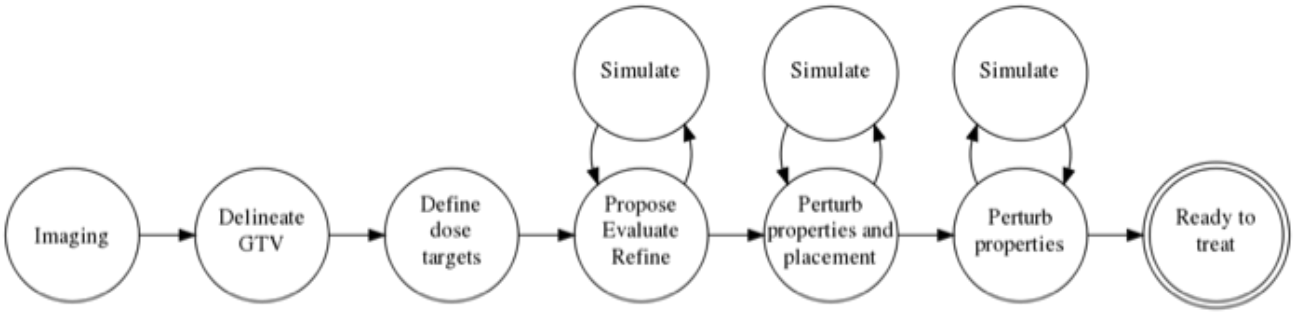

Figure 1. The proposed IPDT pre-treatment workflow. This paper reports our progress through the first four items: moving from medical images to a mesh description, defining targets, proposing a plan, and visualizing the results. The simulation, visualization, and evaluation capabilities reported enable future systematic perturbation to conduct sensitivity analysis and will ultimately support a general algorithmic IPDT planning approach.

by Lund University (with SpectraCure AB)[6] and the University Health Network Toronto[7] have taken this approach, generating dose-volume histograms (DVHs) as the primary method of expressing goals and evaluating plan quality in line with radiation therapy practice.

Due to its ability to deal with a broad range of geometries and material types, Monte Carlo simulation has been widely used as a reference to validate other techniques, but often in simplified geometries (eg. spherical cavity[8] or planar slabs[9]) using symmetry to reduce computing demands.

Aiming to bridge the best of both methods, the geometric flexibility of tetrahedral meshes with the generality of MC methods, we summarize here our progress in implementing the proposed interstitial PDT dose simulation engine, based on our FullMonte 3D tetrahedral Monte Carlo (MC) software. The workflow from medical images through to fluence field visualization is outlined in the next section. Portions of it including the MC simulation kernel and dose-volume histogram generation, are discussed elsewhere. The new infrastructure presented here provides the the necessary platform for plan input, calculation, and result visualization, which will naturally generalize to permit algorithmic plan perturbation, sensitivity analysis, and optimization.

2. WORKFLOW

2.1 Image delineation

In this instance, patient images with mostly pre-existing annotations for organs were provided by the medical physics training program at Princess Margaret Cancer Centre*. The images were contoured using the commercial Pinnacle3 (Philips) radiation treatment planning software by drawing curves on top of a CT DICOM file, with 2mm spacing between slices. Contours are required to delineate both regions of interest (organs at risk and clinical targets), and regions with optical properties significantly different from the surrounding tissue. This differs somewhat from the radiation treatment planning workflow, since in that case the imaging and treatment modalities are similar enough to allow direct imaging of the relevant absorption coefficients.

2.2 Mesh construction

Starting the the input curves, we use the open-source Qhull software[10] to compute a Delaunay tetrahedralization[11] which connects each point to its "nearest" neighbours. From there, we devised an algorithm to identify the tetrahedra within a given region based on the identities of the curves and the geometric relations between slices. The Delaunay tetrahedralization of the curve point set and tetrahedron region assignment produce a mesh which closely reflects the region segmentation of the input curve set but contains a number of low-quality tetrahedra,[12] as shown for a bladder case in Fig 2.

*Thanks are due to Robert Weersink for providing access to the files, and for completing some missing contours

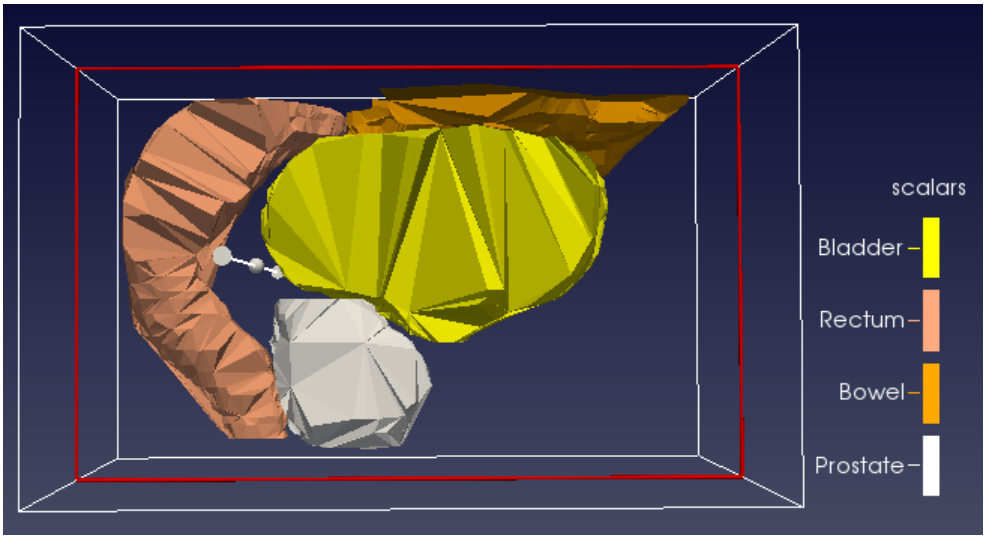

Figure 2. Cut showing the poor-quality, long, thin tetrahedra (108k elements total) before refinement. Tetras composing the surrounding tissue (not shown to avoid clutter) are particularly large. Long tetrahedra cause excess fluence averaging along their long axis, which hurts result quality.

2.3 Mesh refinement

While the mechanics of photonic Monte Carlo simulation are completely different from a finite-element solution of the diffusion approximation, issues of mesh quality which are well-known from finite element analysis also appear in this context. Since the Monte Carlo simulation kernel integrates absorption events over the volume of each element to determine fluence indirectly, the presence of elements which are elongated along a particular axis are particularly detrimental. If the fluence gradient runs along the long axis of the tetrahedron, then it will average out the high fluence at one end with the lower fluence at the other to produce an intermediate value, discarding the gradient information. Fig 2 illustrates the presence of long thin tetrahedra in the pre-refinement mesh. Using the open-source TetGen software[13] created by Si,[14] the mesh is refined using both a maximum-volume criterion and a maximum edge-radius ratio to produce a higher-quality mesh as shown in Fig 3.

Limiting the volume of each element reduces the maximum volume over which averaging can take place, which intuitively increases spatial resolution at the cost of additional complexity. Jointly with the maximum-volume criterion, the edge-radius ratio serves to limit the length of an element (and hence degree of spatial averaging) along any given axis. These involve trade-offs since meshes with larger numbers of tetrahedra take more space to store, more computation to produce, and more time to simulate. The present arrangement is by no means "optimal" and could use further exploration. However, for the meshes shown here the MC run times remain reasonable, and examination of the mesh indicates that fortunately the mesh is at its finest when the fluence is highest by virtue of the density of curve points around the source.

2.4 Plan definition

For this instance, the plan requires optical property assignments, definition of the target tissue, and placement of one or more sources. We have produced an interactive 3D tool which takes the mesh produced above and permits the user to position an isotropic point source or extended line source within the defined treatment volume. Optical properties (μ_s,μ_a,g,n) are specified through a text file. For the bladder case, we used a single isotropic point source located within the bladder over the prostate and near the rectum.

2.5 Monte Carlo simulation

The FullMonte MC kernel has been described previously[15] as the fastest open source optical simulator to our knowledge. Other preceding work[16] has discussed the merits of tetrahedra over voxels, and discussed alternative

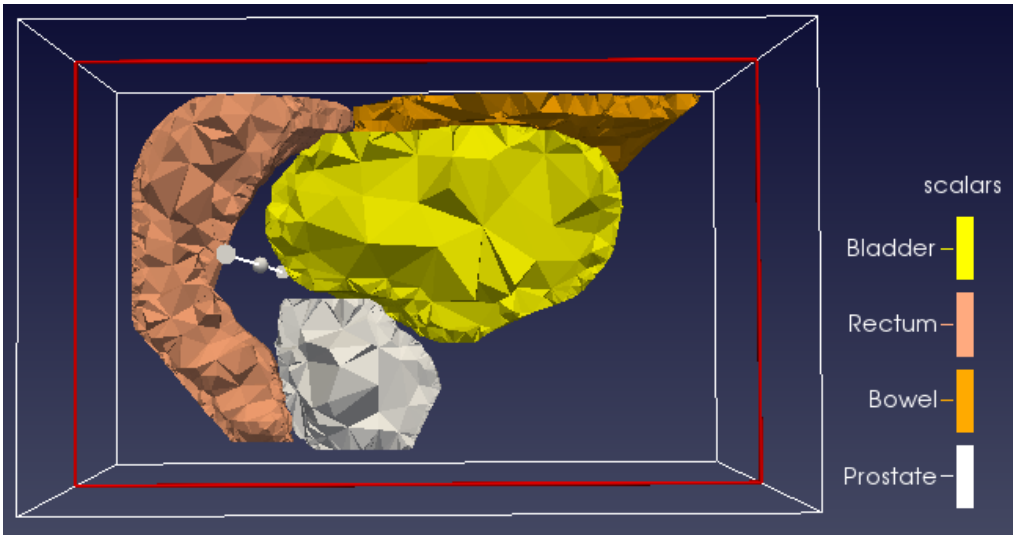

Figure 3. Cut showing the refinement achieved based on tetra volume and shape (1.7 M tetrahedra total). Note the grading of mesh size from fine at the boundaries to coarse in the interior. Propagation calculations are affected only by correct assignment of material to region. Fluence resolution depends on the size and shape of the elements over which absorption is integrated. Particularly in the bladder, the fluence in the interior is not particularly interesting.

geometry formulations.[17] Briefly, our kernel uses the familiar "hop-drop-spin" technique to simulate photon packets as they move through the volume defined by the tetrahedral mesh. The packet weight (energy) deposited in each element is accumulated to yield the fluence distribution as a function of position throughout the mesh. We find MC methods attractive because they are simple, inherently parallel, do not require approximations of the problem physics (unlike diffusion approximation[5] or SP_3 methods[18]), and provide a continuous trade-off between result uncertainty as expressed by statistical variance and run time. Coarse results can be provided quickly, and finer (higher-confidence/lower-variance) results can be generated given time at the option of the user (or algorithm designer) as needed for a given application. Understanding the level of precision required for a given objective is crucial to effective use of the MC method.

2.6 Dose visualization and evaluation

In order to visualize the results of our simulations, we use the Visualization Toolkit (VTK[19]), an open-source project supporting quality 3D visualizations. It permits panning, rotating, slicing, contouring, colouring, and other very useful transformations of the tetrahedral mesh input data, which was used to generate the figures in this publication.

For quantitative evaluation, we recently reported[20] using an open-source mouse model that generation of a dose-volume histogram from a Monte Carlo simulation can be considered a very effective variance reduction scheme. The variance of an individual mesh element depends on its volume, transport coefficient ($\mu_t = \mu_a + \mu_s$), and fluence. Elements which are smaller, contain a material with a longer mean free path, or receive lower fluence generally produce fewer absorption events to average over, leading to higher output variance. However, when many elements are sorted into a dose-volume histogram the variance of that histogram is considerably less than the variance of the constituent elements due to smoothing imposed by sorting. This suggests that even high-spatial-resolution tetrahedral MC methods can be used to generate high-confidence DVHs in seconds, a performance which is competitive with diffusion-approximation methods that offer inferior flexibility in terms of source types and material properties.

3. PRELIMINARY RESULTS: BLADDER CASE

In our recent work,[20] we have demonstrated that the computation run time of Monte Carlo IPDT plan evaluation is not as prohibitive as generally believed when using our highly-optimized software. In fact, for purposes

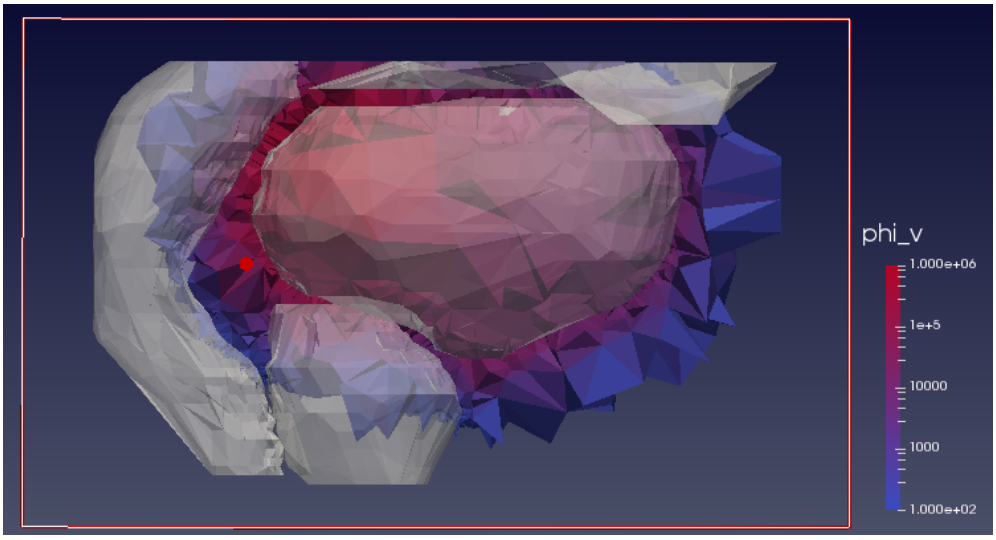

Figure 4. Fluence cut (blue-red gradient) showing log scale with four orders of magnitude in fluence (au). The organ contours are overlaid in transparent white to provide a sense of location. The isotropic point source is located within the bladder, above the prostate and near the rectum (the deepest red tetra). Note that by construction the tetras are smaller where the gradient is larger, while areas remote from treatment (eg. at right) have coarser tetras with much lower fluence. The clipping plane can be moved interactively, and the view can be rotated/zoomed in 3D. This high-quality, high-dynamic-range simulation of 10M packets required 27 seconds to complete on a laptop.

of dose-volume-histogram generation it can be reasonably competitive (around one second) with a diffusion-approximation-based solver on a CPU platform with potential for significant outperformance when using hardware acceleration. Similar to those findings from a mouse model, a few seconds of run time on a laptop generate a stable dose-volume histogram (DVH) using MC methods even on the much larger mesh. This is due at least in part to the strong confinement of light to a small region, such that even though the mesh itself is large (over a million elements) the relevant number of elements is much smaller.

To generate the high-quality and high-dynamic-range data used to generate the fluence cut plane depicted in Fig 4 took slightly longer at 27 seconds. It shows low-noise data over four orders of magnitude, which far exceeds the likely range of interest in IPDT treatment planning. If we consider that the maximum fluence in an overdosed region might be an order of magnitude above the target, conservatively assume that sensitive tissue has a threshold dose an order of magnitude lower than the target, and stipulate further that the proportion of tissue receiving 1/10th of the target dose is required, only about three orders are required so the runtime could be shortened.

4. CONCLUSIONS AND FUTURE WORK

The PDT community realizes the importance of the ability to simulate the fluence fields, both as one of the three essential constituents of photodynamic dose, and as a means of sensing the treatment parameters to provide closed-loop control. The Monte Carlo method is widely acknowledged as the "gold standard" (both in terms of quality and cost) for performing such biophotonic simulations but has yet to be widely applied for IPDT treatment planning and analysis in general 3D geometries. We have demonstrated here that such an approach can yield quality results in a reasonable time.

The other prerequisites for use of tetrahedral MC are of course construction of the input tetrahedral meshes, and making sense of the output data. With our recent progress in generating meshes from 3D medical images and result visualization, we have completed many of the engineering tasks necessary to support the evaluation portion of the FullMonte treatment planning flow outlined in the introduction. By showing a practical workflow from medical image though resulting light field, we believe the present work demonstrates for the first time that

tetrahedral MC is a practical technique for interstitial PDT modeling in human patients. Outstanding items for further development and research to enable full clinical impact are discussed below.

4.1 Meshing

The present meshing process is somewhat cumbersome, requiring user input and trial-and-error to arrive at an acceptable output. We believe it also produces an unnecessarily high number of tetrahedral elements, because of the high density of input curve points required to make the present region-assignment algorithm work. While positive for the spatial resolution of the output, smaller tetrahedra require more memory to store, more computing time to produce a mesh (1-2 minutes for the meshes here), and more time to compute the simulations. A new algorithm based on computing Delaunay surface triangulations is in progress which should reduce or eliminate the burden on the user, as well as give a coarser starting mesh which the user may refine if needed to their desired level of resolution.

4.2 Clinical outcome validation

We seek to obtain images and treatment plans for existing clinical IPDT cases and run prospective simulations to assess the ability of our simulations to predict PDT treatment outcome. Bladder, head-and-neck, and brain cases are all particularly interesting due to their heterogeneous optical properties, anatomical complexity, and the presence of sensitive organs nearby. We believe that Monte Carlo methods will provide new insight into these cases, particularly in the presence of voids with differing refractive indices (eg. head and neck), and low-scattering regions (bladder, CSF in the brain).

4.3 Variability studies

Using the presented platform, we will perform a number of studies to quantify the impact on the fluence field of a number of sources of variation which prove problematic in clinical PDT. For instance, the inter-patient and day-to-day variability of optical properties for a given tissue type may be on the order of 20%, which will significantly alter the fluence distribution within the treatment volume. In some indications, source placement precision may also be an issue due to needle flexion and tissue mechanical properties. Such work will also guide further development of PDT treatment delivery mechanisms by identifying the relative importance of the parameters which determine the probability of delivering a safe and effective light dose.

4.4 Optical monitoring placement

Interstitial PDT can be envisioned for highly heterogeneous anatomical regions, often with limited access for placement of light sources and measuring probes. With large variability as discussed above, verification of the fluence dose delivered is an important concern for clinical PDT: there are significant consequences to both over-treatment of organs at risk, potentially leading to morbidity, and under-treatment permitting continued disease. Given the ability to simulate a broad range of optical properties and placements in actual clinical datasets, we have the opportunity to investigate the question of optimal use of the limited optical measurement opportunities for dose measurement, either directly through photosensitizer phosphorescence , or through measurements of the light fluence. By placing additional monitoring probes at points maximally sensitive to a patient's individual optical properties, we can aim to provide truly personalized PDT treatments through on-line monitoring of optical properties and fluence.

4.5 Optimization

Given a fast, high-quality forward solution and the means to evaluate results both qualitatively and quantitatively, the question of algorithmic optimization naturally arises. We plan to study current practices in radiation therapy with respect to plan figures of merit (objective functions) and non-convex optimization techniques. With appropriate adaptations for the differing mechanisms of action in PDT, we will implement numerical optimization using the simulation and evaluation kernel presented here. Of particular interest for PDT due to its greater variability is the notion of a *robust* planning process which can produce plans with a high probability of successful delivery in spite of variability. Rather than working on point estimates, we envision a process which works stochastically over the range of variability in error sources to identify and attempt to mitigate the impact of variability.

4.6 Computational acceleration

Despite several successes accelerating layered[21] and voxelized simulations[22] using Graphics Processing Units (GPUs), we are not aware of any high-speed tetrahedral-based MC code for GPU. We continue to investigate that possibility, but currently believe that it would be difficult to achieve high performance due to the random memory access patterns required by the tetrahedral mesh. Even if performance gains are possible, such a system is likely to be quite demanding of both power and cooling which makes it less attractive in a clinical setting.

Taking another approach following on Lo's work accelerating modelling of planar geometries,[23] we have already demonstrated an FPGA hardware implementation with a limited (48k-element) mesh size[24] which achieves significant performance gains (4x, scalable to 16x) and massive gains in performance-per-watt (67x) over our CPU code. Efforts are currently on-going to realize the system using the IBM OpenPOWER CAPI (Coherent Accelerator-Processor Interface) platform, a unique interface which will enable the FPGA to access the host's memory permitting essentially unlimited mesh size while retaining the power and performance advantages. On completion, such a system could explore potentially dozens of treatment plans per second, or generate extremely high-confidence fluence maps and parameter sensitivities in seconds - all within a small form factor and power budget.

ACKNOWLEDGMENTS

Robert Weersink provided the contoured anatomy and valuable discussion regarding radiation treatment planning. The visualizations presented here rely extensively on the open-source Visualization Toolkit (VTK) and Paraview, both produced by Kitware Inc. The authors also acknowledge funding from Altera Corporation (JC/VB), IBM (JC/VB), the Canadian Natural Sciences and Engineering Research Council (VB), the Canadian Institutes of Health Research via a Banting and Best Doctoral Research Award (JC), and the Ontario Ministry of Health and Long-Term Care (LL).

REFERENCES

1. T. C. Zhu and J. C. Finlay, "Prostate PDT Dosimetry," *Photodiagnosis and Photodynamic Therapy* **3**, pp. 234–236, 2006.
2. B. C. Wilson and M. S. Patterson, "The physics, biophysics and technology of photodynamic therapy," *Physics in medicine and biology* **53**, pp. R61–109, May 2008.
3. M. A. Biel, "Photodynamic Therapy," *Methods in Molecular Biology* **635**, pp. 281–293, 2010.
4. J. Cassidy, V. Betz, and L. Lilge, "Towards PDT Treatment Planning," in *37th Annual Meeting of the American Society for Photobiology*, American Society for Photobiology, (San Diego, CA), 2014.
5. S. L. Jacques and B. W. Pogue, "Tutorial on diffuse light transport," *Journal of biomedical optics* **13**(4), p. 041302, 2008.
6. J. Swartling, J. Axelsson, G. Ahlgren, K. M. Kälkner, S. Nilsson, S. Svanberg, K. Svanberg, and S. Andersson-Engels, "System for interstitial photodynamic therapy with online dosimetry: first clinical experiences of prostate cancer.," *Journal of biomedical optics* **15**(5), p. 058003, 2010.
7. S. R. H. Davidson, R. a. Weersink, M. a. Haider, M. R. Gertner, A. Bogaards, D. Giewercer, A. Scherz, M. D. Sherar, M. Elhilali, J. L. Chin, J. Trachtenberg, and B. C. Wilson, "Treatment planning and dose analysis for interstitial photodynamic therapy of prostate cancer.," *Physics in medicine and biology* **54**, pp. 2293–313, Apr. 2009.
8. J. L. Sandell, T. C. Zhu, and J. C. Finlay, "A study of light fluence rate distribution for PDT using MC simulation," in *Optical Methods for Tumor Treatment and Detection: Mechanisms and Techniques in Photodynamic Therapy XX*, D. H. Kessel and T. Hasan, eds., **7886**, pp. 788619:1–6, SPIE, (San Francisco, CA), Feb. 2011.
9. L. Wang, S. L. Jacques, and L. Zheng, "MCML - Monte Carlo modeling of light transport in multi-layered tissues," *Computer Methods and Programs in Biomedicine* **47**(2), pp. 131–146, 1995.
10. C. B. Barber, D. P. Dobkin, and H. Huhdanpaa, "The Quickhull Algorithm for Convex Hulls," *ACM Trans. Math. Softw.* **22**, pp. 469–483, Dec. 1996.
11. Joseph O'Rourke, *Computational Geometry in C*, Cambridge University Press, 1998.

12. S.-W. Chen, T. K. Dey, and J. R. Shewchuk, *Delaunay Mesh Generation*, CRC Press, Boca Raton, FL, 2012.
13. H. Si, "TetGen," 2013.
14. H. Si, "TetGen: A quality tetrahedral mesh generator and three-dimensional delaunay triangulator," *Weierstrass Institute for Applied Analysis and Stochastic, Berlin, Germany* , 2006.
15. J. Cassidy, L. Lilge, and V. Betz, "FullMonte: a framework for high-performance Monte Carlo simulation of light through turbid media with complex geometry," in *Proc SPIE BiOS*, **8592**, pp. 85920H–14, SPIE, (San Francisco, CA), Feb. 2013.
16. H. Shen and G. Wang, "A study on tetrahedron-based inhomogeneous Monte Carlo optical simulation.," *Biomedical optics express* **2**, pp. 44–57, Jan. 2010.
17. Q. Fang, "Mesh-based Monte Carlo method using fast ray-tracing in Plücker coordinates.," *Biomedical optics express* **1**, pp. 165–75, Aug. 2010.
18. A. D. Klose and E. W. Larsen, "Light transport in biological tissue based on the simplified spherical harmonics equations," *Journal of Computational Physics* **220**, pp. 441–470, Dec. 2006.
19. W. J. Schroeder, L. S. Avila, and W. Hoffman, "Visualizing with VTK: a tutorial," *Computer Graphics and Applications, IEEE* **20**, pp. 20–27, Sept. 2000.
20. J. Cassidy, V. Betz, and L. Lilge, "Treatment plan evaluation for interstitial photodynamic therapy in a mouse model by Monte Carlo simulation with FullMonte," *Frontiers in Physics* **3**(6), 2015.
21. E. Alerstam, W. C. Y. Lo, T. D. Han, J. Rose, S. Andersson-Engels, and L. Lilge, "Next-generation acceleration and code optimization for light transport in turbid media using GPUs," *Biomedical Optics Express* **1**(2), pp. 658–675, 2010.
22. Q. Fang and D. a. Boas, "Monte Carlo simulation of photon migration in 3D turbid media accelerated by graphics processing units.," *Optics express* **17**, pp. 20178–90, Oct. 2009.
23. W. C. Y. Lo, K. Redmond, J. Luu, P. Chow, J. Rose, and L. Lilge, "Hardware acceleration of a Monte Carlo simulation for photodynamic therapy treatment planning," *Journal of Biomedical Optics* **14**(1), p. 014019, 2009.
24. J. Cassidy, L. Lilge, and V. Betz, "Fast , Power-Efficient Biophotonic Simulations for Cancer Treatment using FPGAs," in *International Symposium on Field-Configurable Custom Computing Machines*, IEEE, (Boston, MA), 2014.