

Computational Modeling and Real-Time Control of Patient-Specific Laser Treatment of Cancer

D. F UENTES,¹ J. T. ODEN,¹ K. R. D ILLER,² J. D. HAZLE,³ A. ELLIOTT,³ A. SHETTY,³ and R. J. STAFFORD³

¹Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, TX 78712, USA;

²Department of Biomedical Engineering, The University of Texas at Austin, Austin, TX 78712, USA; and ³Department of Imaging Physics, University of Texas M.D. Anderson Cancer Center, Houston, TX 77030, USA

(Received 10 June 2008; accepted 22 December 2008; published online 16 January 2009)

Abstract—An adaptive feedback control system is presented which employs a computational model of bioheat transfer in living tissue to guide, in real-time, laser treatments of prostate cancer monitored by magnetic resonance thermal imaging. The system is built on what can be referred to as cyberinfrastructure, a complex structure of high-speed network, large-scale parallel computing devices, laser optics, imaging, visualizations, inverse-analysis algorithms, mesh generation, and control systems that guide laser therapy to optimally control the ablation of cancerous tissue. The computational system has been successfully tested *in vivo*, canine prostate. Over the course of an 18 min laser-induced thermal therapy performed at M.D. Anderson Cancer Center (MDACC) in Houston, Texas, the computational models were calibrated to intra-operative real-time thermal imaging treatment data and the calibrated models controlled the bioheat transfer to within 5 °C of the predetermined treatment plan. The computational arena is in Austin, Texas and managed at the Institute for Computational Engineering and Sciences (ICES). The system is designed to control the bioheat transfer remotely while simultaneously providing real-time remote visualization of the on-going treatment. Post-operative histology of the canine prostate reveal that the damage region was within the targeted 1.2 cm diameter treatment objective.

Keywords—Hyperthermia, Real-time computing, Medical imaging, Cancer treatment, Cyberinfrastructure, PDE constrained optimization.

INTRODUCTION

The use of modern methods of computer modeling and simulation to guide and predict the outcome of medical procedures is a development of monumental importance that could eventually change and enrich all

of medical science and significantly enhance the health and quality of life of the human species. One recent development in computational science that can greatly increase the “fidelity and predictive power of computer simulations is the use of feedback control systems that integrate computer models with “old measurement systems so as to allow near-real-time calibration of the models. Such systems are collectively called Dynamic Data Driven Application Systems (DDDAS),⁶ and, in theory, such systems can be used to deliver predictive simulations of unprecedented accuracy and reliability.

This paper describes a canonical feedback control system developed to control and predict the outcome of laser therapies for prostate cancer. The system involves the integration of a long list of technologies and new methodologies, including magnetic resonance thermal imaging (MRTI), computer visualization, laser optics, high-speed networks, nonlinear dynamic bioheat transfer models of heterogeneous tissue, adaptive meshing, high-performance parallel computing, cell-damage and heat-shock protein models, inverse analysis, calibration, model validation, signal processing, optimal control algorithms, and error estimation and control. In our particular DDDAS, the fact that these diverse technologies and systems are connected across a high-speed computational grid connecting remote sites, an imaging and medical unit in Houston and a modeling and computational unit in Austin, qualifies, in our opinion, as an example of cyberinfrastructure.³

The basic hypothesis underlying the creation of this system is that cancer cells can be eradicated when subjected to sufficient heat for a sufficient period of time. Such heat sources can be supplied by inserting optical “fibers into the infected tissue and then supplying energy by means of a laser. If one can predict the evolution of temperature “fields generated by such heat sources, then the location of the “fiber and the frequency and power of the laser can be controlled to

Address correspondence to D. Fuentes, Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, TX 78712, USA. Electronic mail: fuentes@ices.utexas.edu, oden@ices.utexas.edu, kdiller@mail.utexas.edu, jhazle@di.mdacc.tmc.edu, Andrew.Elliott@di.mdacc.tmc.edu, Anil.Shetty@di.mdacc.tmc.edu and jstafford@di.mdacc.tmc.edu; <http://dddas.ices.utexas.edu>

minimize damage to healthy tissue while maximizing damage to tumor cells. The development of calibrated models that accurately simulate bioheat transfer in prostate tissue is thus a key component of the system. But, as can be anticipated, there are many complicating factors in developing such a control system, including, for example, the development of computational models of cell damage and heat-shock protein expression. This research aims to focus on minimally invasive thermal therapy delivery designed to provide local treatment of focal disease, such as those designed to replace, or act as an adjuvant to, conventional surgery. The real-time monitoring of temperature using MRTI is useful in these treatments in order to increase both efficacy and safety due to the rapid heating and high temperatures involved. The proposed technology further increases this efficacy by helping to optimize the conformality of the treatment as well as minimize risks associated with damaging nearby critical structures.

The core simulation tool embeds experimental real-time thermal imaging data within the nonlinear Pennes bioheat transfer model. The thermal imaging data are provided by an MRTI device that follows the actual therapy at the imaging laboratory at M.D. Anderson Cancer Center (MDACC) in Houston, Texas. The MRTI technology developed over the past decade is a modification of existing MRI technology to use fast imaging sequences to acquire larger imaging volumes in the same time with comparable temperature sensitivity and to provide a 3D temperature field as a function of position and time in the living tissue. The optimal control system is guided by simulations performed at the computational modeling arena with the Institute for Computational Engineering and Sciences (ICES) in Austin, Texas with high-bandwidth network connections to computers at the Texas Advanced Computing Center (TACC). As noted, through accurate computer prediction, the location and power of the laser can, in principle, be modulated to maximize damage to the tumor while minimizing damage to healthy tissue. Computer prediction may be used to accurately control the affected tissue response through a collection of imaging based measurements about how the complex physiological system is responding to the surgery and to make changes in the treatment plan based on an intelligent understanding of the physiological pathways and feedback mechanisms to affect the surgical outcome. This represents a significant extension of modern simulation methods as a means for real-time calibration and validation of models are made possible in this system. The possibility of simulating bioheat transfer in cancerous tissues of mice was established in the work of Rylander et al.^{13,17} that was a precursor of the present study. In these works,

primitive models of cell damage and heat-shock-protein expression were also proposed and calibrated by means of in vitro experiments.

Following this introduction, we describe the computational infrastructure constructed to implement the control strategy, including the assembly of various software modules and the new software units developed to carry out various functions for the systems. In the **Real-Time Optimization** section, we describe inverse analysis and optimization algorithms developed around a computational model based on a nonlinear parabolic partial differential equation that is a generalization of the classical Pennes model.¹⁵ The control paradigm is based on the derivation of backward-in-time adjoint problems associated with various cost functionals. It is shown that a single mathematical and computational framework can be developed for adjoints corresponding to control cell damage, HSP expression, and numerical accuracy in approximations of specific quantities of interest. The computational model of heat transfer is constructed using adaptive hp-nite elements⁷ and a Crank-Nicolson scheme for approximating the temporal behavior of the temperature field. The **Results** section contains discussions of results of actual physical trials employing the full system for predicting thermal response and controlling laser input. This involves so-called phantom tests, run on an agar gel in which a canine prostate is implanted, and MRTI-monitored in vivo experiments, guided by the control system, involving laser treatment on prostates of living canines. The results clearly establish that the system can provide accurate depictions of in vivo thermal histories in canines and that the computer simulations provide adequate information for effectively controlling laser inputs and locations. Major conclusions of the study are collected in the **Conclusions and Future Directions** section.

CYBERINFRASTRUCTURE AND WORK FLOW

We begin by describing the so-called cyberinfrastructure developed to implement a DDDAS. This involves developing a computational grid which connects the experimental arena, at which the laser treatment for prostate cancer is done, and the computational arena, at which the computer modeling, visualizations, and high performance computing is carried out. The former is built around the MRTI laboratory at M. D. Anderson Cancer Center in Houston, Texas and the latter is managed at ICES in Austin, Texas, approximately 150 miles away. Communications between these arenas are handled via batch sftp over a commercial GigE Internet connection. The full system is depicted symbolically in Fig.1.

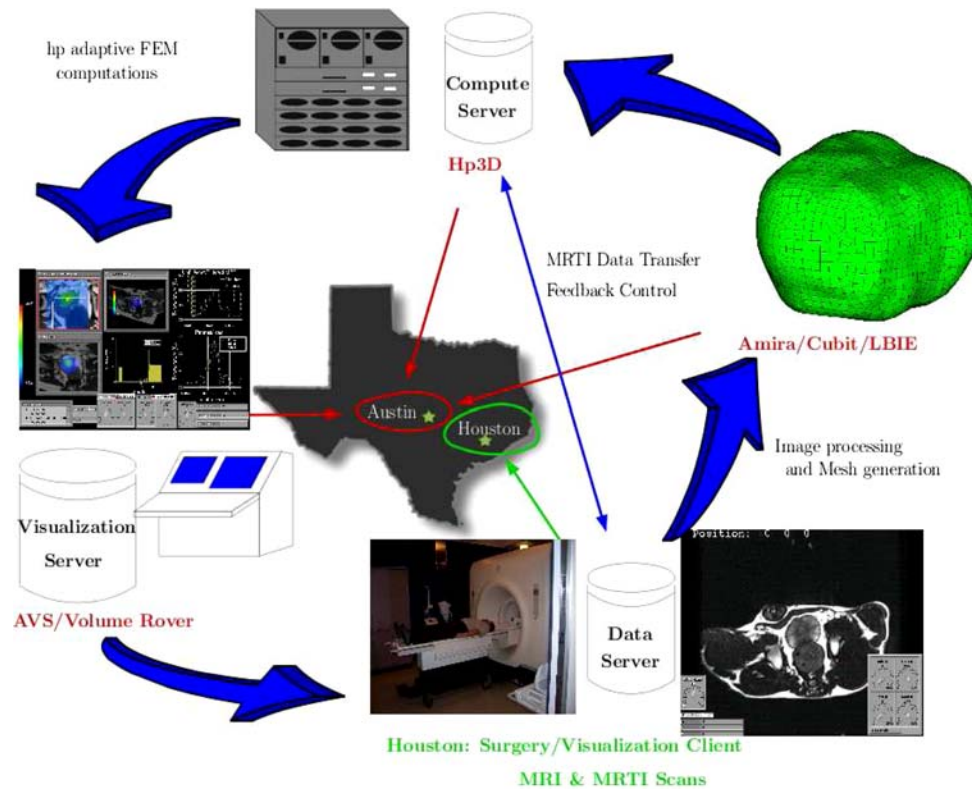


FIGURE 1. The communication architecture used to control the laser treatment process over a distance of 150 miles is shown. The continual interaction of the computational models, implemented at the Texas Advanced Computing Center in Austin, with the thermal imaging data, acquired at M.D. Anderson Cancer Center in Houston, provides the feedback control. Visualization of the treatment is rendered in Austin and provided to Houston, remotely.

The principal software modules and the data flow are depicted in Fig. 2.

The control strategy involves multiple steps. Several days prior to the laser treatment, the prostate area of the therapy subject, in our case a laboratory canine, is scanned using a 1.5 T clinical MR scanner (Excite HD, GEHT, Waukesha, WI, USA). This pre-operative, patient-specific data are transmitted to the computational arena where an initial 3D finite element mesh approximating the prostate geometry is generated. Several mesh generation packages are available, including the Level Set Boundary-Interior-Exterior (LBIE) Mesher¹ presented.⁹ A pipeline, Fig. 3, of commercial software may also be used. AMIRA, <http://www.amiravis.com>, is used to perform a semi-automatic segmentation and create a faceted surface representation of the prostate. CUBIT, cubit.sandia.gov, is then used to create a pillowed hexahedral mesh and apply boundary conditions. The majority of the time in this pipeline is spent in the semi-automatic segmentation of the prostate. An expert user must identify and extract the geometry of the prostate

from the pre-operative images. Prior to treatment, reservations for access to the supercomputers as well as the MRI scanner are made, and initial optimal laser parameters are identified such as the location of the endpoint of the optical fiber, laser power as a function of time, etc. A trial simulation of the therapy is performed using initial bioheat transfer data, some taken from handbooks,¹⁰ to determine the initial location of the laser fiber. The initial parameters are ultimately corrected during the calibration phase of the process using temperature data generated via MRTI. This is described fully in **Results** section. The entire system, computation, data transfer, visualization, etc., is tested repeatedly to ensure stable code execution during the treatment window. The laser treatment may take place over a period of up to 18 min, but shorter procedures may also be adequate.

We employ an FDA cleared system for MR temperature monitoring of laser-induced thermal therapy (Visualase[®], BioTex Inc., Houston, TX, USA) modified to be controlled from remote supercomputers in Austin. The Visualase uses a diffusing interstitial laser fiber to provide the heat source to heat the tissue to an ablative temperature regime. This system is monitored remotely at the computational arena in Austin. Upon

¹Software available at: <http://cvcweb.ices.utexas.edu/cvc>

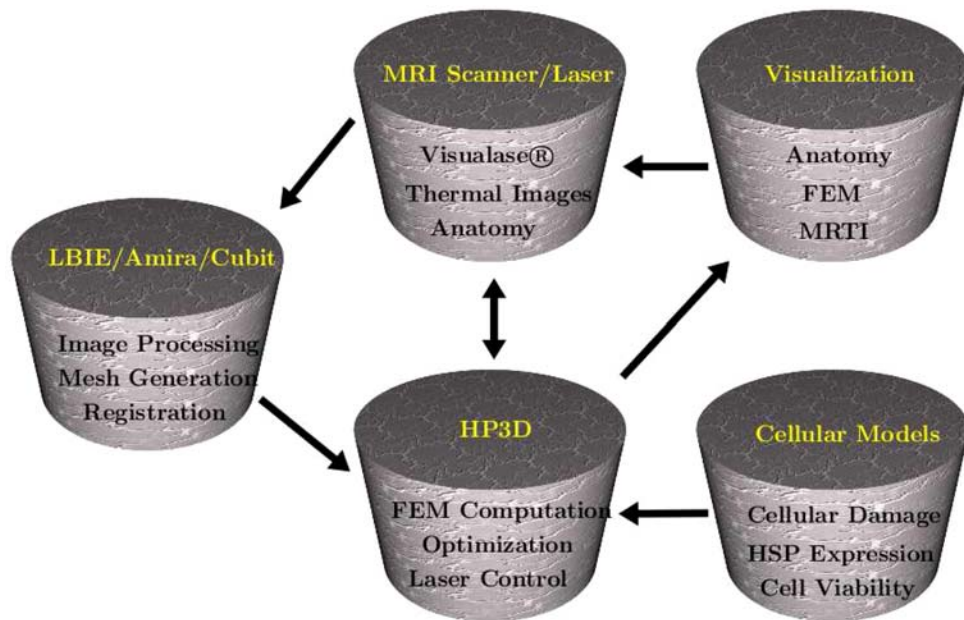


FIGURE 2. There are five main software modules which together represent the computational infrastructure. A suite of programs handles the anatomy and real-time thermal image acquisition; the geometry extraction and finite element mesh generation; the simultaneous real-time visualization of the FEM, anatomical, and thermal image data; the implementation of the cellular damage models; and finally the PDE constrained optimization. HP3d uses the Visualase software to control the laser source. The arrows indicating the data flow between the control system modules is illustrated.




FIGURE 3. AMIRA is used for 3D semi-automatic segmentation in the *XY*, *XZ*, and *YZ* planes of the anatomical data. The data are output as a facet file. The output facet file is then input into CUBIT to create a pillowed hexahedral mesh.

initiation of the treatment process, the software processes the intra-operative MRI data to register the computational domain and real-time MRTI data drives the calibrations aligning the parameters of the Pennes bioheat transfer model to the patient's biological tissue values. As new thermal imaging data are acquired intermittently, the computational prediction is compared to the measurements of the real-time thermal images and the differences seen are used to update the computations of the optimal laser parameters needed to achieve the original planned treatment. If the model is incorrectly calibrated and the thermal images begin to detect unexpected high temperatures, the laser power is turned off. The model must then be recalibrated before proceeding. The entire system is built around "ve main computing modules as indicated in Fig. 2.

The treatment protocol is broken into four stages, as illustrated in Fig. 4. The problems of real-time calibration, optimal control, and goal-oriented error




FIGURE 4. The treatment process is divided into four stages. (1) MRTI thermal image data is acquired and used for model calibration; (2) the time span of actual calibration computations; (3) the time span for optimal temperature/damage/HSP computations; and (4) the optimal laser control parameters are applied to the biological domain. In the event of the detection of an unexpectedly high temperature within the biological domain, a fail-safe signal shuts off the laser power.

estimation are solved in parallel by separate groups of processors working together as indicated in Fig. 5. Before any thermal image transfer occurs, an initial prediction of the entire time history of temperature/

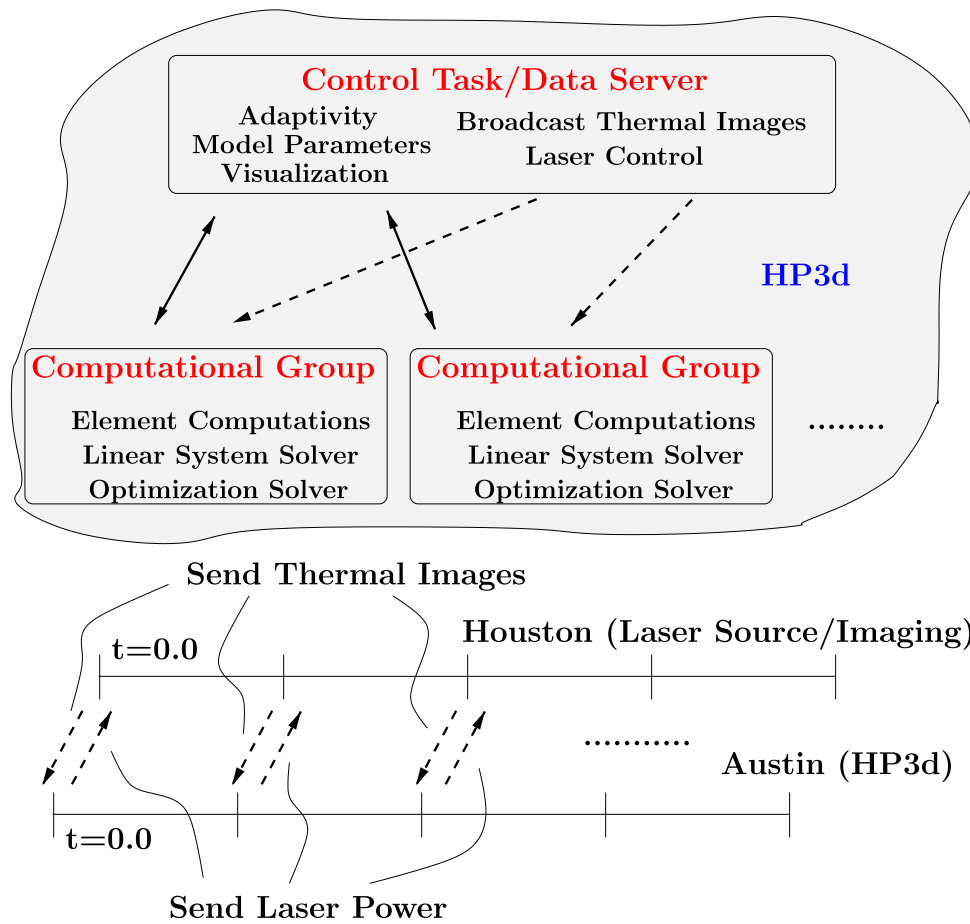


FIGURE 5. The finite element computations are performed on a parallel computing architecture using multiple groups of processors to simultaneously solve the disjoint PDE constrained optimization problems of the control system. A single processor acts as a Control Task/Data Server to orchestrate the entire treatment. The processor is used to gather and broadcast the individual solutions of the computational groups and broadcast filtered thermal images to individual computational groups as requested. The imaging implicitly controls the laser power output. As a new thermal image is acquired by HP3d in Austin, the estimated power wattage for the next time interval is transmitted to the laser and is physically generated during the therapy.

damage/HSP from textbook or pre-calibrated material properties is computed. The time window of the prediction is the entire length of the treatment $[0, t_f]$. During the “rst stage, the biological domain is pulse heated during time $[0, t_0]$ and MRTI thermal image data is acquired from time $[0, t_1 > t_0]$ for the heating as well as the cooling. The images from $[0, t_1]$ are used to calibrate the bioheat transfer model. The time window $[t_1, t_2]$ is allotted to “nish the calibration computation. The calibration computations are check-pointed intermittently. The timing of the optimal control computations are implicitly based on the acquisition of the thermal image at t_{im2} . At time t_2 , a second group of processors uses the most recent checkpoint of the calibration parameter update and proceeds to optimize the power from $[t_3, t_f]$ ($t_3 > t_2$); the time window $[t_2, t_3]$ is provided to “nish the optimal control computations. At time t_3 , the optimizer is stopped and the “nal prediction is made over the time window $[t_3, t_f]$. During this fourth and “nal stage, the optimal laser control

parameters modulate the laser power for the remainder of the treatment. In the event of the detection of an unexpectedly high temperature within the biological domain, a fail-safe shuts off the laser power.

A Control Task unit orchestrates the treatment while simultaneously acting as a data server of the thermal images. Periodically during treatment, the group-wise optimization solutions and error estimates are collected on the Control Task data repository. A skeleton of the entire FEM mesh of the biological domain is stored on the Control Task. Using the collective error estimates, a mesh re“nement strategy is computed on the Control Task and both the collective optimization solutions and re“nement strategy are broadcast to the individual Computational Groups. The Control Task/Data Server, Fig. 5, reads in the thermal images from disk, “lters the thermal images to remove noise, broadcasts the thermal images to each computational group as needed, and transmits the laser power to Houston. The timing of the laser power control is implicit through the image

acquisition by the HP3d, Fig. 5, in Austin. As new thermal images are written to disk physically in Houston, they are transferred to disk at the computing center (TACC) in Austin. When the Data Server detects that the full set of thermal images for a time instance is available, the predicted power to be used for the next time interval is sent to Houston.

A high-level abstraction of the HP3d^{8,2} hp-adaptive “nite element solver is given in Fig. 5. The models are implemented efficiently in parallel to meet the demands of rapid calibration and rapid optimal control calculations. The “nite element solver is implemented at the Austin site and run at the computing center (TACC) through a network connecting ICES to the center. The software infrastructure is built from the PETSc⁴ parallel computing paradigm and the TAO⁵ parallel optimization library. AVS¹ is used in conjunction with a VNC server for remote visualization. AVS coroutines are used to manage and coordinate the simultaneous visualization of the MRI anatomical image, MRTI thermal image, and “nite element data sets. From a computational point of view, the orchestration of a successful laser treatment must handle the problems of registration, calibration, and optimal control invisibly to the surgeon, and provide a priori visualization of the outcome of the treatment for use as guidance for the actual treatment.

REAL-TIME OPTIMIZATION

We employ a nonlinear generalization of the classical Pennes⁵ model of bioheat transfer as the base model for simulating the evaluation of temperature in the biological domain due to energy supplied by the laser. The domain, Ω , shown in Fig. 6, is heated by a laser heat source within the time interval $[0, \tau]$. The entire time




FIGURE 6. The equations of bioheat transfer are derived from the conservation of energy applied to an arbitrary sub-volume. The biological domain is composed of a cancerous region (red) and a region of healthy tissue. A laser source located at x_0 heats the domain within the time interval $[0, \tau]$.

single continuous medium that may have spatially varying properties and mass flux does not occur across the boundary. The heat transfer between blood and tissue is the defining characteristic of biological heat transfer; as blood perfuses the complex vasculature networks embedded in tissue, it may act as a significant spatially varying heat source or sink. The biological domain Ω is composed of a region of healthy tissue Ω_H and a region of cancerous tissue Ω_C , $\Omega = \Omega_H \cup \Omega_C$.

The nonlinear Pennes bioheat transfer model is given as follows:

$$\begin{aligned} \rho c_p \frac{\partial u}{\partial t} - \nabla \cdot (k(u, x, \beta) \nabla u) + \omega(u, x, \beta) c_{\text{blood}} (u - u_a) &= Q_{\text{laser}}(\beta, x, t) \quad \text{in } \Omega \\ Q_{\text{laser}}(\beta, x, t) &= 3P(t) \mu_a \mu_{\text{tr}} \frac{\exp(-\mu_{\text{eff}} \|x - x_0\|)}{4\pi \|x - x_0\|^2} \quad \mu_{\text{tr}} = \mu_a + \mu_s (1 - g) \\ \mu_{\text{eff}} &= \sqrt{3\mu_a \mu_{\text{tr}}} \\ -k(u, x, \beta) \nabla u \cdot n &= h(u - u_{\infty}) \quad \text{on } \partial\Omega_C \\ -k(u, x, \beta) \nabla u \cdot n &= \mathcal{G} \quad \text{on } \partial\Omega_N \\ u(x, 0) &= u^0 \quad \text{in } \Omega \end{aligned}$$

span of the laser treatment occurs within the interval $[0, T]$, $[0, \tau] \subset [0, T]$. The Pennes model describes bioheat transfer as the conservation of energy applied to a motionless nondeforming mass of human tissue. The muscle, fat, blood, connective tissue is represented as a

The initial conditions, u^0 , are taken as the baseline body temperature. The density of the continuum is denoted ρ and the specific heat of blood is denoted c_{blood} [J/kg K]. On the Cauchy boundary, $\partial\Omega_C$, h is the coefficient of cooling and u_{∞} is the ambient temperature. The prescribed heat flux on the Neumann boundary, $\partial\Omega_N$, is denoted \mathcal{G} . The optical...thermal response to the laser source, $Q_{\text{laser}}(\beta, x, t)$, is modeled as the classical

²Software available at <http://dddas.ices.utexas.edu>

spherically symmetric isotropic solution to the transport equation of light within a laser-irradiated tissue.²⁰ β denotes an array of bioheat transfer model parameters.

$$\beta \equiv (k_0(x), k_1, \tilde{k}_3, \hat{k}_3, \omega_0(x), \omega_1, \tilde{\omega}_3, \hat{\omega}_3, P(t), \mu_a, \mu_s, x_0)$$

$P(t)$ is the laser power as a function of time μ_a and μ_s are laser coefficients related to laser wavelength and give probability of absorption and scattering of photons, respectively. The anisotropic factor is denoted μ_a and x_0 denotes the position of laser photon source. The scalar-valued coefficient of thermal conductivity, k [J/s m K], is a function of the local temperature, u , and is assumed to be of the form (Fig. 7)

$$k(u, x) = k_0(x) + k_1 \operatorname{atan}(\tilde{k}_2(u - \hat{k}_3))$$

where $k_0(x)$ [J/s m K], k_1 [J/s m K], \tilde{k}_2 [1/K], \hat{k}_3 [K] $\in \mathbb{R}$. The function $k_0(x)$ is allowed to vary over the spatial dimension to capture the biological tissue heterogeneity. The media are thus assumed to be thermally isotropic (referring to heat flux) but spatially heterogeneous. The perfusion coefficient is given by ω [kg/s m³]. Empirical evidence suggests that [kg/s m³], is a smooth, monotone increasing bounded function of temperature (Fig. 7). We use the characterization,

$$\omega(u, x) = \omega_0(x) + \omega_1 \operatorname{atan}(\tilde{\omega}_2(u - \hat{\omega}_3))$$

where ω_0 [kg/s m³], ω_1 [kg/s m³], $\tilde{\omega}_2$ [1/K], $\hat{\omega}_3$ [K] $\in \mathbb{R}$. Note that $\omega_0(x)$ is allowed to vary over the spatial dimension as the blood perfusion within the necrotic core of a cancerous tumor or the blood perfusion within a damaged tissue is expected to be significantly lower than the surrounding healthy tissue.

The main problems of interest is the real-time solution of the calibration of the model coefficients and the optimal control of the laser. The problem of adaptive mesh refinement is presented in Odeh et al.¹⁴ The model parameter available for control of the laser

heat source is the power as a piece-wise constant function of time, $P(t)$. The material properties available for model calibration are the thermal conductivity parameters $k_0(x), k_1, \tilde{k}_3, \hat{k}_3$ the perfusion parameters $\omega_0(x), \omega_1, \tilde{\omega}_3, \hat{\omega}_3$ and the coefficients of scattering and absorption, μ_s and μ_a . The position of the laser tip, $x_0 = (x_0, y_0, z_0)$, may be used in both calibration and optimal control. Notice the explicit dependence of thermal conductivity, $k(u, x, \beta)$, perfusion, $\omega(u, x, \beta)$, and laser source, $Q_{\text{laser}}(\beta, x, t)$, on β . The temperature field resulting from the application of the heat source is an implicit function of the model parameters.

$$u(\beta, x, t) \quad t \in [0, T] \quad x \in \Omega \quad \beta \in \mathbb{P}$$

where \mathbb{P} is defined in (5).

Calibration of Pennes Model

Given a fixed laser power, $P(t)$, the problem of model calibration is to find the set thermal conductivity parameters, $k_0(x), k_1, \tilde{k}_3, \hat{k}_3$, blood perfusion parameters, $\omega_0(x), \omega_1, \tilde{\omega}_3, \hat{\omega}_3$, optical coefficients, μ_a, μ_s , and laser parameters x_0 , that minimize the $L_2(0, T; L_2(\Omega))$ norm of the difference between the predicted temperature field and an ideal temperature field, $u^{\text{ideal}}(x, t)$.

$$\begin{aligned} Q(u(\beta, x, t), \beta) &= \frac{1}{2} \|u(\beta, x, t) - u^{\text{ideal}}(x, t)\|_{L_2([0, T]; L_2(\Omega))}^2 \\ &= \frac{1}{2} \int_{\Omega} \int_0^T (u(\beta, x, t) - u^{\text{ideal}}(x, t))^2 dt dx \end{aligned} \quad (1)$$

$dx = dx_1 dx_2 dx_3$ being a volume element. We describe here a calibrated deterministic model; stochasticity is to be addressed in future work. The ideal field for the calibration problem is the experimentally determined temperature field, $u^{\text{exp}}(x, t)$, obtained from in vivo MRTI

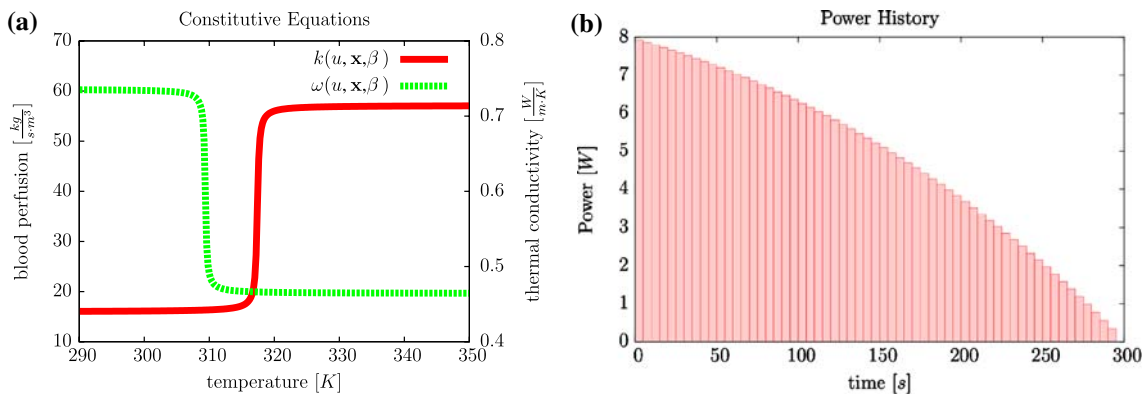


FIGURE 7. (a) Typical temperature dependencies of the blood perfusion and thermal conductivity material data used in the Pennes model are plotted against the left and right axis, respectively. (b) The power is piecewise constant in time. The duration of an individual power pulse is determined by the image acquisition time.



FIGURE 8. The 1.5-T clinical scanner used in this research is shown. The average data set sizes obtained from this scanner are 3.7 MB ($256 \times 256 \times 30$ voxel), 3D anatomical data for planning, and .32 MB ($256 \times 256 \times 5$ voxel), thermal images every 4–6 s for monitoring progression of treatment.

data. The thermal imaging data, as well as the anatomical data, is acquired using a 1.5 T MRI scanner (Sigma Excite HD, Milwaukee, WI, USA), Fig. 8, using a 2D temperature sensitive echo planar imaging sequence; the scanner is located at the experimental (MRTI) arena in Houston. The MRTI process itself involves a phase subtraction technique. Each consecutive phase image is subtracted from a baseline image that was acquired prior to heating.¹⁹ These phase difference images are proportional to temperature difference, Δu . To convert phase to temperature other factors including echo time, magnetic field strength, gyro magnetic ratio water ppm shift with temperature are required. The total temperature change is given as the cumulative sum of the individual temperature differences. The image acquisition time of the MRTI data is 4...6 s. Given a file size of approximately .32 MB for one time instance of thermal imaging data and average bandwidth transferring data from Houston to Austin of approximately .2 MB/s, the time to transfer the thermal imaging data is roughly 1 s.

Optimal Control of the Treatment

The optimal treatment of a focal disease is defined by the following criteria:

- ... Maintain the functionality of the healthy region of the tissue.
- ... Maximize damage to the cancerous region.

Various metrics based on temperature, damage, or HSP expression are used to evaluate the degree to which the criteria is satisfied. For optimal control, the laser power, $P(t)$, is varied as a function of time to minimize the objective function and the material parameters, $k(u, x, \beta)$, $\omega(u, x, \beta)$, μ_{sr} and μ_a are fixed. An optimal

laser position x_0 may be pre-computed. However, the surgeon's ability to achieve the pre-computed position is limited and the actual treatment position of the laser may need to be calibrated from the thermal imaging data.

Temperature-Based Optimal Control

Temperature-based optimal control is similar to the calibration problem, the goal being to find the set of model parameters that minimize the space-time norm of the difference between the computed temperature field $u(\beta, x, t)$ and an ideal field $u^{ideal}(x, t)$. The mathematical structure is the same as in (1). However, the ideal field for temperature-based optimal control maximizes damage to cancerous tissue while minimizing damage to healthy tissue,

$$u^{ideal} = \begin{cases} 37^\circ\text{C} & x \in \Omega_H \\ 60^\circ\text{C} & x \in \Omega_C \end{cases}$$

Damage-Based Control

Damage-based control attempts to minimize the L_2 -norm of the difference between an ideal damage field and the computed damage field,

$$\begin{aligned} Q(u(\beta, x, t), \beta) &= \frac{1}{2} \|\varphi(u(\beta, x, t)) - D^{ideal}(x)\|_{L_2(\Omega)}^2 \\ &= \frac{1}{2} \|D(x) - D^{ideal}(x)\|_{L_2(\Omega)}^2 \end{aligned}$$

An ideal damage field, $D^{ideal}(x)$, maximizes damage to the cancerous region while maintaining functionality of the healthy region. Within the cancerous region there is a threshold damage value, denoted D_C^{ideal} , above which all cancerous cells will die. Similarly, within the healthy region there is a threshold damage value, denoted D_H^{ideal} , below which the functionality is maintained for all healthy cells.

$$D^{ideal}(x) \equiv \begin{cases} D_C^{ideal} & x \in \Omega_C \\ D_H^{ideal} & x \in \Omega_H \end{cases}$$

The functional that acts on the temperature field and returns the corresponding damage field is denoted ϕ ,

$$\begin{aligned} \phi : L_2(0, T, H^1(\Omega)) &\rightarrow L_2(\Omega) \\ \phi(u(\beta, x, t)) &= D(x) \quad \forall x \in \Omega \end{aligned}$$

The damage field, $D(x)$, reflects the time-temperature thermal injury resulting from the application of the heat source.

An Arrhenius model and a Two-State model¹² are used to quantify the damage from the thermal assault on the biological domain, Ω . The Arrhenius model is a classical empirical law used in early models of thermal effects and properties of molecular averages. For the Arrhenius model, ϕ represents the damage index and for the Two-State model, ϕ represents the cell viability.

$$\varphi(u(\beta, x, t)) = D(x) = \begin{cases} \int_0^T A e^{\frac{-E_a}{Rt}} dt & \text{Arrhenius} \\ \int_0^T \frac{e^{-(\frac{h}{u} + \alpha t + \beta)}}{1 + e^{-(\frac{h}{u} + \alpha t + \beta)}} dt & \text{Two-State} \end{cases} \quad (2)$$

The constants E_a , A , and R are known constants of the Arrhenius model and h , α , and β are known constants of the Two-State model determined from in vitro experiments on cell samples.¹² Note that the numerical values of D_C^{ideal} and D_H^{ideal} for the Arrhenius model will be different from D_C^{ideal} and D_H^{ideal} for the Two-State model.

Heat Shock Protein (HSP) Based Control

Heat shock protein (HSP) based control is similar to damage control. It is assumed that the functional mapping the temperature “field to the corresponding HSP expression “field is known:

$$H: L_2(0, T, H^1(\Omega)) \rightarrow L_2(\Omega) \quad H(u(\beta, x, t))(x) \in L_2(\Omega)$$

The constitutive data mapping the temperature “field to the HSP expression “field is obtained from in vitro cellular data.¹⁶ An ideal HSP expression “field, $H^{ideal}(x)$, maximizes the protective effect of heat shock proteins within healthy tissue while minimizing heat shock protein expression within the cancerous region.

A summary of the objective functions is provided as follows:

$$Q(u(\beta, x, t), \beta) = \begin{cases} \frac{1}{2} \int_{\Omega} \int_0^T (u(x, t) - u^{ideal}(x))^2 dt dx & \text{calibration/temp. based} \\ \frac{1}{2} \int_{\Omega} \left(\int_0^T A e^{\frac{-E_a}{Rt}} dt - D^{ideal}(x) \right)^2 dx & \text{Arrhenius} \\ \frac{1}{2} \int_{\Omega} \left(\int_0^T \frac{e^{-(\frac{h}{u} + \alpha t + \beta)}}{1 + e^{-(\frac{h}{u} + \alpha t + \beta)}} dt - D^{ideal}(x) \right)^2 dx & \text{Two-State} \\ \frac{1}{2} \int_{\Omega} (H(x) - H^{ideal}(x))^2 dx & \text{HSP} \end{cases} \quad (3)$$

PDE Constrained Optimization

The mathematical structure of the calibration and optimal control problems is the same: Find the set of model parameters β^* , that minimizes a given objective function, Q , over a parameter manifold, \mathbb{P} ,

$$\text{Find } \beta^* \in \mathbb{P} \text{ s.t.} \\ Q(u(\beta^*), \beta^*) = \inf_{\beta \in \mathbb{P}} Q(u(\beta), \beta)$$

subject to the constraints of a variational form of the Pennes PDE, $C(u, \beta; v): \mathcal{V} \times \mathbb{P} \times \mathcal{V} \rightarrow \mathbb{R}$

$$\forall v \quad C(u, \beta; v) = 0 \\ = \int_0^T \int_{\Omega} \rho c_p \frac{\partial u}{\partial t} v \, dx dt \\ + \int_0^T \int_{\Omega} k(u, x, \beta) \nabla u \cdot \nabla v \, dx dt \\ + \int_0^T \int_{\Omega} \omega(u, x, \beta) c_{blood} (u - u_a) v \, dx dt \\ + \int_0^T \int_{\partial \Omega_N} g v \, dA dt \\ - \int_0^T \int_{\Omega} Q_{laser}(\beta, x, t) v \, dx dt \\ + \int_0^T \int_{\partial \Omega_C} h(u - u_{\infty}) v \, dx dt \quad (4)$$

The variational form (4) is nonlinear with respect to the “rst two arguments, u, β , but linear with respect to the test function, v . The proper mathematical setting for the trial and test functions is the space of mappings from the time domain to the appropriate function space, $H^1(\Omega)$,

$$u(\beta, x, t) \equiv (u(\beta, t))(x): (0, T) \rightarrow H^1(\Omega)$$

such that the H^1 norm of the mapping, u , is L_2 integrable in time,

$$\int_0^T \|u(\beta, \cdot, t)\|_{H^1}^2 dt < \infty$$

The appropriate space containing such mappings is $\mathcal{V} \equiv L_2(0, T; H^1(\Omega))$. A well-defined solution to the variational problem (4) may be shown to exist for model parameters, β , belonging to a parameter space \mathbb{P} ,

$$\mathbb{P} = \left\{ \beta \in L_{\infty}(\Omega) \times \mathbb{R}^3 \times L_{\infty}(\Omega) \times \mathbb{R}^3 \times L_{\infty}([0, T]) \times \mathbb{R}^5 : \right. \\ \left. \begin{aligned} 0 < k_* < k_0(x) + k_1 \tan(k_2(u - k_3)) < k^* < \infty \\ 0 < \omega_* < \omega_0(x) + \omega_1 \tan(\omega_2(u - \omega_3)) < \omega^* < \infty \end{aligned} \right\} \quad (5)$$

over which the operators are coercive, bounded, and type- \mathcal{M} .¹⁸

The gradient of the objective function with respect to the control variables is the work horse of our optimization schemes. The objective function is guaranteed to decrease in the direction of the gradient and is also used to approximate the Hessian in a quasi-Newton scheme. A Newton scheme for optimization is not considered in this work. The real-time nature of the control system does not permit the efficient computation of the Hessian of the objective function. An adjoint solution method of the optimization problems is used to compute the gradient. The adjoint method considered employs framework of differential calculus in Banach spaces and clearly demonstrates how the use of an adjoint variable eliminates the need to solve for sensitivities of the solution. An analytical expression for gradient of the objective function with respect to the control variables is determined under the assumption that the objective function and temperature “eld are Fréchet-differentiable.³

Considering the Pennes model solution, u , as a function of the control parameters, β , and the objective function, Q , as a function of the Pennes model solution, u , and the control parameters, β , the composition of the objective function with the Pennes model solution, denoted as $Q \diamond u$, may be considered as a Fréchet-differentiable mapping solely from parameter space to \mathbb{R} .

$$\begin{aligned} u(\beta): \mathbb{P} &\rightarrow \mathcal{V} \\ Q(u, \beta): \mathcal{V} \times \mathbb{P} &\rightarrow \mathbb{R} \\ Q \diamond u \equiv Q(u(\beta), \beta): \mathbb{P} &\rightarrow \mathbb{R} \end{aligned}$$

The derivative of the composition, $Q \diamond u$, is of interest and is computed under the following “rst order approximations of the solution and objective function.

$$\begin{aligned} u(\beta + \hat{\beta}) &= u(\beta) + u'(\beta)\hat{\beta} \\ Q(u + \hat{u}, \beta + \hat{\beta}) &= Q(u, \beta) + D_u Q(u, \beta)\hat{u} + D_\beta Q(u, \beta)\hat{\beta} \end{aligned}$$

Here $u': \mathbb{P} \rightarrow \mathcal{V}$ denotes the Fréchet derivative of the solution with respect to the control variables, $D_u Q: \mathcal{V} \rightarrow \mathbb{R}$ is the partial Fréchet derivative with respect to the solution, and $D_\beta Q: \mathbb{P} \rightarrow \mathbb{R}$ is the partial Fréchet derivative with respect to the control variables.

For functions defined on the Cartesian products of spaces, the notion of partial Fréchet differentiability is understood as follows:

Partial Fréchet-differentiability: Let X denote the direct sum of the space of trial functions and parameter space, $X = \mathcal{V} \oplus \mathbb{P}$. For U open, $U \subset X$, and $F: U \rightarrow Y$ consider the point $x = (u, \beta) \in X$. Define

$$\begin{aligned} f_u(z): \mathcal{V} &\rightarrow \mathbb{R} & f_u(z) &= F(z, \beta) \\ f_\beta(s): \mathbb{P} &\rightarrow \mathbb{R} & f_\beta(s) &= F(u, s) \end{aligned}$$

The Fréchet derivative of $f_u(z)$ at $z = u$ and $f_\beta(s)$ at $s = \beta$ are the partial derivatives F in the direction of u and β ; respectively.

$$\begin{aligned} D_u F(x): \mathcal{V} &\rightarrow \mathbb{R} & D_u F(x) &= Df_u(u) \\ D_\beta F(x): \mathbb{P} &\rightarrow \mathbb{R} & D_\beta F(x) &= Df_\beta(\beta) \end{aligned}$$

The Fréchet derivative of F at (u, β) is the sum of its partial derivatives.²

The chain rule holds for a composition of Fréchet differentiable functions²; the derivative of the composition is computed as follows,

$$\begin{aligned} D(Q \diamond u): \mathbb{P} &\rightarrow \mathbb{R} \\ D(Q \diamond u) &= D_u Q(u, \beta)u'(\beta) + D_\beta Q(u, \beta) \end{aligned} \quad (6)$$

Notice that to compute the gradient of the objective function, Q , the sensitivities, u' , must be computed. The variation of the PDE constraints with respect to the model parameters may be used to compute the sensitivities. Only model parameters for which a solution to the governing PDE exists are of interest (i.e. $\beta: C(u(\beta), \beta; v) = 0 \forall v$). Within this hyper-space the variation is identically zero.

$$\begin{aligned} (D_u C u' + D_\beta C): \mathbb{P} &\rightarrow \mathbb{R} \\ D_u C(u, \beta; v)u'(\beta) + D_\beta C(u, \beta; v) &= 0 \quad \forall v \in \mathcal{V} \end{aligned} \quad (7)$$

where $D_u C: \mathcal{V} \rightarrow \mathbb{R}$ is the partial Fréchet derivative of the PDE with respect to the solution and $D_\beta C: \mathbb{P} \rightarrow \mathbb{R}$ is the partial Fréchet derivative of the PDE with respect to the model parameters. The adjoint method of computing the gradient of the objective function may now be formulated.

Theorem. Given $\beta \in \mathbb{P}$, the derivative of the objective function is

$$D(Q \diamond u) = D_\beta Q(u(\beta), \beta) - D_\beta C(u(\beta), \beta; p) \quad (8)$$

where $u \in \mathcal{V}$ is the solution of the Pennes model and $p \in \mathcal{V}$ is the solution to the adjoint problem

$$\begin{aligned} \text{Find } p \in \mathcal{V}: \\ \langle D_u C(u(\beta), \beta; p), \hat{u} \rangle_{\mathcal{V} \times \mathcal{V}} \\ = \langle D_u Q(u(\beta), \beta), \hat{u} \rangle_{\mathcal{V} \times \mathcal{V}} \quad \forall \hat{u} \in \mathcal{V} \end{aligned}$$

Proof. For a “xed β , the solution to the adjoint problem, p , satisfies

$$\langle D_u C(u(\beta), \beta; p), u' \rangle_{\mathcal{V} \times \mathcal{V}} = \langle D_u Q(u(\beta), \beta), u' \rangle_{\mathcal{V} \times \mathcal{V}}$$

From the “rst variation of the PDE constraints (7), we have that

³The notion of Fréchet differentiability is covered in Atkins et al.²

$$D_u C(u(\beta), \beta; p) u'(\beta) = -D_\beta C(u(\beta), \beta; p)$$

The gradient follows from a direct substitution in (6) to eliminate the sensitivities:

$$\begin{aligned} D(Q \diamond u) &= D_u Q(u, \beta) u'(\beta) + D_\beta Q(u, \beta) \\ &= D_\beta Q(u, \beta) - D_\beta C(u(\beta), \beta; p) \end{aligned}$$

□

Recall that $D(Q \diamond u)$ is an operator on the parameter space, \mathbb{P} , and the desired solution, $\hat{\beta}$, is such that

$$D(Q \diamond u)(\hat{\beta}) = 0 \iff \langle D(Q \diamond u)(\hat{\beta}), \hat{\beta} \rangle = 0 \quad \forall \hat{\beta} \in \mathbb{P}$$

As a concrete example, the strong form of the adjoint problem is derived from the Fréchet derivatives used to derive (8). The following Taylor series expansions are used in computing the derivatives of the variational form of Pennes model:

$$\begin{aligned} k(u + \hat{u}, \beta + \hat{\beta}) &= k(u, \beta) + \hat{u} \frac{\partial k}{\partial u} + \hat{k}_0 \frac{\partial k}{\partial k_0} \\ &\quad + \hat{k}_1 \frac{\partial k}{\partial k_1} + \hat{k}_2 \frac{\partial k}{\partial k_2} + \hat{k}_3 \frac{\partial k}{\partial k_3} + \text{h.o.t.} \\ \omega(u + \hat{u}, \beta + \hat{\beta}) &= \omega(u, \beta) + \hat{u} \frac{\partial \omega}{\partial u} + \hat{\omega}_0 \frac{\partial \omega}{\partial \omega_0} \\ &\quad + \hat{\omega}_1 \frac{\partial \omega}{\partial \omega_1} + \hat{\omega}_2 \frac{\partial \omega}{\partial \omega_2} + \hat{\omega}_3 \frac{\partial \omega}{\partial \omega_3} + \text{h.o.t.} \end{aligned}$$

An alternative formulation of (4) must be used to account for the initial condition. The initial condition is accounted for in the variational form by integrating by parts on the time derivative.

$$\begin{aligned} &\int_{\Omega} \rho c_p [u(x, 0) p(x, 0) - u(x, T) p(x, T)] dx \\ &= \int_0^T \int_{\Omega} \rho c_p \left(\frac{\partial u}{\partial t} p + \frac{\partial p}{\partial t} u \right) dx dt \end{aligned}$$

Recognizing that $u^0(x) = u(x, 0)$ and not testing the equation at the final time, $p(x, T) = 0$, leads to the following

$$\begin{aligned} C(u, \beta; p) &= - \int_0^T \int_{\Omega} \rho c_p \frac{\partial p}{\partial t} u dx dt \\ &\quad + \int_0^T \int_{\Omega} k(u, \beta) \nabla u \cdot \nabla p dx dt \\ &\quad + \int_0^T \int_{\Omega} c_{\text{blood}} \omega(u, \beta) (u - u_a) p dx dt \\ &\quad + \int_0^T \int_{\partial \Omega} h u p dAdt + \int_{\Omega} u^0(x) p(x, 0) dx \end{aligned}$$

Applying the definition of the partial Fréchet derivative to the variational PDE gives

$$\begin{aligned} C(u + \hat{u}, \beta; p) &= - \int_0^T \int_{\Omega} \rho c_p \frac{\partial p}{\partial t} (u + \hat{u}) dx dt \\ &\quad + k(u + \hat{u}, \beta) \nabla (u + \hat{u}) \cdot \nabla p dx dt \\ &\quad + \int_0^T \int_{\Omega} c_{\text{blood}} \omega(u + \hat{u}, \beta) ((u + \hat{u}) - u_a) p dx dt \\ &\quad + \int_0^T \int_{\partial \Omega} h (u + \hat{u}) p dAdt + \int_{\Omega} u^0(x) p(x, 0) dx \end{aligned}$$

Substituting the Taylor series expansion for k and ω and rearranging terms gives

$$C(u + \hat{u}, \beta; p) - C(u, \beta; p) = \left[\begin{aligned} & - \int_0^T \int_{\Omega} \rho c_p \frac{\partial p}{\partial t} \hat{u} dx dt + \int_0^T \int_{\partial \Omega} h \hat{u} p dAdt \\ & + \int_0^T \int_{\Omega} \left(k(u, \beta) \nabla \hat{u} \cdot \nabla p + \hat{u} \frac{\partial k}{\partial u} \nabla u \cdot \nabla p \right) dx dt \\ & + \int_0^T \int_{\Omega} \left(\omega(u, \beta) \hat{u} p + \hat{u} \frac{\partial \omega}{\partial u} (u - u_a) p \right) dx dt \\ & + \int_0^T \int_{\Omega} \left(\hat{u} k(u, \beta) \nabla \hat{u} \cdot \nabla p + \frac{1}{2} \hat{u}^2 \frac{\partial^2 k}{\partial u^2} \nabla u \cdot \nabla p \right) dx dt \\ & + \dots \end{aligned} \right]$$

Notice that the terms involving the initial condition cancel out. Letting

$$\begin{aligned} D_u C(u, \beta; p) \hat{u} &= \int_0^T \int_{\Omega} \rho c_p \frac{\partial \hat{u}}{\partial t} p \, dx dt + \int_0^T \int_{\partial \Omega} h \hat{u} p \, dA dt \\ &+ \int_0^T \int_{\Omega} k(u, \beta) \nabla \hat{u} \cdot \nabla p + \hat{u} \frac{\partial k}{\partial u} \nabla u \cdot \nabla p \, dx dt \\ &+ \int_0^T \int_{\Omega} \omega(u, \beta) \hat{u} p + \hat{u} \frac{\partial \omega}{\partial u} (u - u_a) p \, dx dt \end{aligned}$$

implies that

$$|C(u + \hat{u}, \beta; p) - C(u, \beta; p) - D_u C(u, \beta; p) \hat{u}| \leq \tilde{C} \|\hat{u}\|^2 \quad (9)$$

where the higher-order terms of the thermal conductivity, $\partial^2 k / \partial u^2$, and perfusion, $\partial^2 \omega / \partial u^2$, are assumed bounded and hidden in the constant \tilde{C} . For a particular test/adjoint function, p , the constant, \tilde{C} , also depends on the norm of the test/adjoint function, $\|p\|$; however, because $p \in \mathcal{V}$, the norm is always bounded and (9) holds. It follows from the definition that $D_u C(u, \beta; p)$ is the Fréchet derivative of C . The strong form of the adjoint formulation is as follows.

$$\begin{aligned} & - \rho c_p \frac{\partial p}{\partial t} - \nabla \cdot (k(u, x, \beta) \nabla p) \\ & + \frac{\partial k}{\partial u} (u, x, \beta) \nabla u \cdot \nabla p + \omega(u, x, \beta) p \\ & + \frac{\partial \omega}{\partial u} (u, x, \beta) p (u - u_a) = \delta Q \quad \text{in } \Omega \\ & - k(u, x, \beta) \nabla p \cdot n = h p \quad \text{on } \partial \Omega \\ & p(x, T) = 0 \quad \text{in } \Omega \end{aligned}$$

Notice that the adjoint problem for the calibration, optimal control, and goal oriented error estimates are all essentially the same, the only difference being the source term δQ which depends on the quantity of interest.

Discretization of Equations

The Galerkin representations of the temperature field and adjoint variable are assumed as follows

$$u(x, t) = \sum_{k=1}^{N_{\text{step}}} \sum_{j=1}^{N_{\text{dof}}} \alpha_j^k(t) \phi_j(x) \quad p(x, t) = \sum_{k=1}^{N_{\text{step}}} \sum_{i=1}^{N_{\text{dof}}} \lambda_i^k(t) \phi_i(x)$$

where N_{step} is the number of time steps, N_{dof} is the number of Galerkin coefficients, and ϕ_i 's are the finite element shape functions.

$$\alpha_j^k(t) = \begin{cases} \frac{t_k - t}{t_k - t_{k-1}} \alpha_j^{k-1} + \frac{t - t_{k-1}}{t_k - t_{k-1}} \alpha_j^k, & t \in [t_{k-1}, t_k) \\ 0, & \text{otherwise} \end{cases}$$

$$\lambda_i^k(t) = \begin{cases} \lambda_i^k, & t \in [t_{k-1}, t_k) \\ 0, & \text{otherwise} \end{cases}$$

The time discretization of the power is assumed piecewise constant in time:

$$P(t) = \begin{cases} P_k, & t \in [t_{k-1}, t_k) \\ 0, & \text{otherwise} \end{cases}$$

The spatial variation of the parameters "elds is assumed to have the following Galerkin representation,

$$k_0(x) = \sum_j k_0^j \psi^j(x) \quad \omega_0(x) = \sum_j \omega_0^j \psi^j(x)$$

where $\psi(x)$ are piecewise constant across elements. The test function are assumed piecewise constant in time

$$v(x, t) = \sum_{k=1}^{N_{\text{step}}} \sum_{i=1}^{N_{\text{dof}}} v_i^k(t) \phi_i(x) \quad v_i^k(t) = \begin{cases} v_i^k, & t \in [t_{k-1}, t_k) \\ 0, & \text{otherwise} \end{cases}$$

The governing equations (4) are solved using a Crank-Nicolson time stepping scheme.

The Adjoint Gradient of the quantity of interest is constructed from the derivative of the discretized equations with respect to an arbitrary single model variable, β_i . The gradient of the quantity of interest is computed using the chain rule for derivatives taken with respect to the numerically computed quantity of interest. The initial condition does not depend on the model parameters, $\frac{\partial u_0}{\partial \beta_i} = 0$,

$$\frac{\partial}{\partial \beta_i} Q(u(\beta, x, t), \beta) = \sum_{k=1}^{N_{\text{step}}} \frac{\partial Q}{\partial u_k} \frac{\partial u_k}{\partial \beta_i}$$

As shown in Demkowicz et al.,⁹ the numerical derivative of the quantity of interest with respect to a model variable may be computed as follows.

$$\begin{aligned} & \frac{\partial Q(u, \beta)}{\partial \beta_i} \\ &= \sum_{k=1}^{N_{\text{step}}} \left(\begin{aligned} & -\Delta t_k \int_{\Omega} \frac{\partial k}{\partial \beta_i} (u_{k-\frac{1}{2}}, x, \beta) \nabla u_{k-\frac{1}{2}} \cdot \nabla p_k \, dx \\ & -\Delta t_k \int_{\Omega} c_{\text{blood}} \frac{\partial \omega}{\partial \beta_i} (u_{k-\frac{1}{2}}, x, \beta) (u_{k-\frac{1}{2}} - u_a) p_k \, dx \\ & +\Delta t_k \int_{\Omega} \frac{\partial Q_{\text{laser}}}{\partial \beta_i} (\beta, x, t_k) p_k \, dx \end{aligned} \right) \end{aligned}$$

RESULTS

In this section, results from phantom testing and in vivo experiments of the control system are presented

along with post-treatment histology for comparison with computed predictions. Performance results used in the design of the treatment protocols are also provided. All computations use the GMRES linear solver and block Jacobi preconditioner available from PETSc.⁴ A bound constrained quasi-Newton linesearch method available in TAO⁵ is used for the optimization solution. The linesearch is based on the Wolfe conditions, sufficient decrease and curvature, and uses a quadratic/cubic interpolation scheme for the step length selection.

Phantom Validation Tests

Developmental stages uses a series of so-called phantom experiments constructed to provide a preliminary validation setting for the mathematical and computational models developed in this study; the phantoms are constructed from ex vivo canine prostate tissue embedded within a 1% agar material (Fig. 9). The phantom provides an animal-free method of testing and debugging every aspect of the control system in a controlled and repeatable manner without interference from confounding conditions such as time varying perfusion as expected in their *in vivo* system. The phantom is meshed, registered, and the computational models in Austin are used to control the heating. The heating of the phantom is visible in the MRTI images. Thermal images of the heating are acquired and sent to the HP3d program; HP3d calibrates and optimizes the model parameters, and finally, a visualization of the entire process is provided in Houston. Figure 10 shows a post-calibration time instance of the predicted temperature of the phantom heated with an interstitial laser fiber. A visualization of the anatomy, predicted

temperature contours, predicted damage contours, cutlines of the baseline thermal images, and cutlines of the heating predicted by Pennes model are shown. As there is no perfusion within the excised tissue, Pennes model reduces to a nonlinear heat transfer equation. The anatomy images are used to obtain the initial coordinates of the laser tip in the FEM calculations. As illustrated by the power profile, Fig. 10e, the phantom was pulsed heated with the laser source at 8 W applied power for 30 s. Imaging of the heating and cooling is used for model calibration; in this particular trial, thermal imaging data were only acquired for calibration. The initial pulse used for model calibration is seen to mildly damage the prostate before the optimal control portion of the experiment (Fig. 10b). The optimal power profile shown was computed to simulate damage-based control using the Arrhenius model (2) to optimally heat a tumor of radius 1.5 cm. The profile shown maximizes damage to the tumor and minimizes damage to the surrounding healthy tissue by following a pulse heating with a tapering of the laser power; the computed profile begins at image step id 120 shown in Fig. 10e. The tapering allows heat to diffuse through the tumor without overheating the healthy tissue. The contour plots of the thermal images and the calibrated Pennes model prediction is seen to differ (Fig. 10a and d). As expected from modeling the medium as homogeneous with isotropic heating, the Pennes model prediction shows circular contours compared to the more elliptical shape seen in the thermal image contours. However, the particular time instance illustrates very good outline agreement between the thermal images and the Pennes model prediction (Fig. 10f). Also shown are cutlines of the filtered and unfiltered thermal images (Fig. 10c).

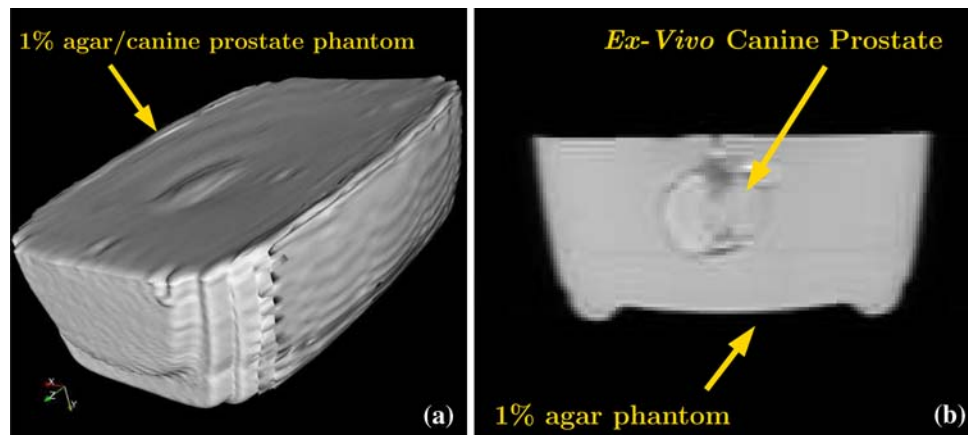


FIGURE 9. The phantom geometry used in the development of the control system is shown. The phantom consists of a canine prostate embedded within a 1% agar gel. (a) Isosurface visualization of MRI images of the geometry of the phantom material. (b) A particular image within the 3D data set shows the *ex-vivo* canine prostate embedded within the 1% agar gel.

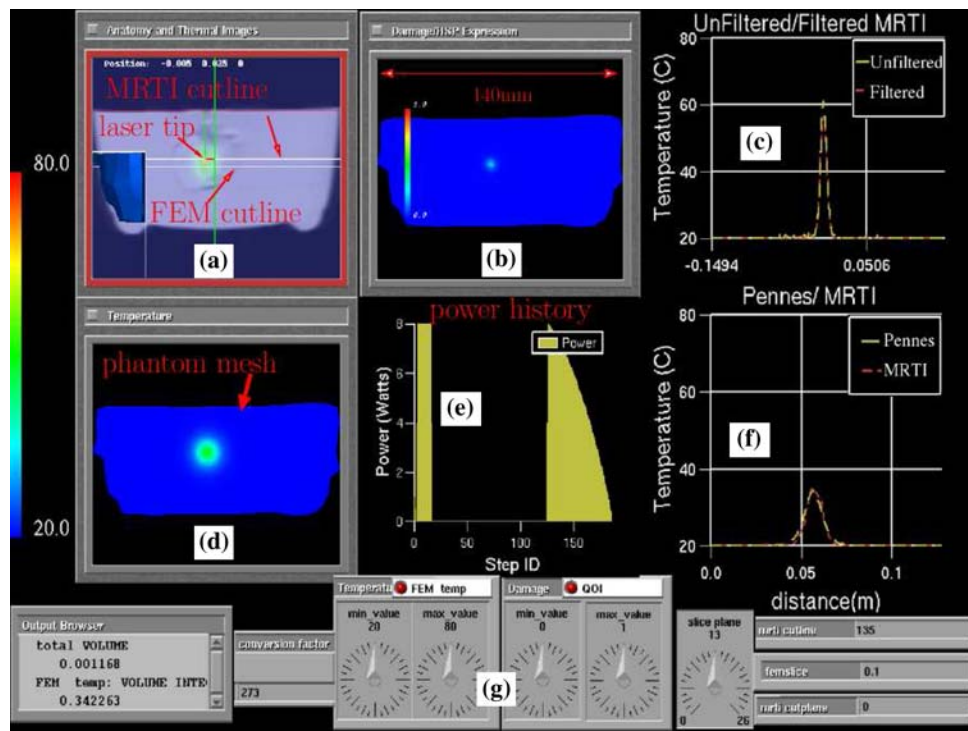


FIGURE 10. Every aspect of the control system has been tested on a phantom constructed of an excised canine prostate embedded within a 1% agar gel. Depicted is the visualization provided during the laser treatment. The color scale, shown at the left, is from 20–80 °C and applies to the color scale in (a) and (d). The power history as a function of image number is shown in (e). Time is implicit in the image number as the image acquisition time is 5 s. As shown, the phantom was heated with a calibration pulse of 8 W for 30 s. The visualization depicts the calibrated model at image number 11. (a) The thermal images are overlaid on the anatomical images; also shown are the cutlines through the thermal images and FEM predictions used in (c) and (f), respectively. The cutlines shown are at different spatial positions. The thermal imaging cutline comparing the filtered, red, and unfiltered, yellow, data are shown in (c); the distance along this cutline is given in meters and the coordinates are with respect to the DICOM images from the scanner. A cutline through the finite element data comparing the Pennes model prediction, yellow, to the thermal imaging data, red, projected onto the finite element mesh is shown in (f); the distance along this cutline is given in meters and the coordinates are with respect to the initial point of the cutline. A contour plot of the FEM temperature prediction of the phantom is shown in (d); the cutplane corresponds to the cutplane shown in (a). (b) An Arrhenius damage model prediction is shown post-calibration pulse. The calibration pulse is seen to cause mild damage. (g) A few of numerous controls needed to precisely display the remote visualization of the treatment are shown. The controls shown may be used to adjust the color scale of the temperature and the slice plane of the thermal/anatomical images as well as the position of the contour slice of the FEM mesh.

Treatment Protocol Design

The execution time and speedup of a representative 10 s bioheat transfer simulation is presented in Fig 11. The execution times represent 10 nonlinear state solves of the Pennes model (10 s time steps) combined with 10 linear adjoint solves to be used in the gradient computation for calibration, optimal control, and/or error estimates. A “nite element mesh consisting of approximately 10,000 degrees of freedom is used to characterize the computational performance. The results presented may be used as a template for extrapolating to more realistic computational models with larger degrees of freedom. Computations were done at the computational arena at TACC on a Dual-Core Linux Cluster. Each node of the cluster contains two Xeon Intel Duo-Core 64-bit processors (4 cores in all) on a single board, as an SMP unit. The core frequency is 2.66 GHz and supports 4 floating-point operations

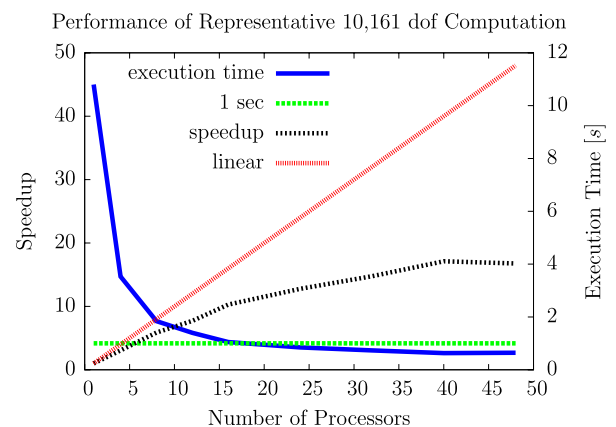


FIGURE 11. Shown is the speedup as a function of processors and execution time for a representative 10 s simulation (10 nonlinear state solve combined with 10 linear adjoint solve on 10000 dof system). The fastest execution time is of the highest priority, the efficiency of the computation is sacrificed for execution time.




FIGURE 12. Pre-operative and intra-operative imaging data of canine prostate and the neighboring anatomy. The canine is laying supine with legs upward. The pre-operative data were used to create a 3D FEM mesh of the prostate consisting of 23,303 dof and 20,064 hexahedral elements. A cropped section of the FEM mesh is shown. The images are reconstructed to 256×256 voxels over a $240 \text{ mm} \times 240 \text{ mm}$ field of view with 36 and 12 acquired slices (3 mm thick with 1.5 mm spacing) being used for the pre-operative planning and intra-operative updating, respectively. The intra-operative image shown was used to locate the interstitial laser fiber within the DICOM coordinate system and update the treatment plan.

per clock period. The fastest time recorded is .67 s. Furthermore, as shown in Bensor et al.,¹¹ the optimizer is seen to be capable of converging to a solution in under 10 function/gradient evaluations. Hence, an estimated time of 2.5 min of computation time is needed for the calibration computations using 150 s worth of data—30 images with a 5-s acquisition time, using a ratio of 10 s of computation time to 1 s of wall time, and allowing for 10 function/gradient evaluations.

In Vivo Experiments

The control system, in its entirety, has been twice tested on in vivo canine prostate tissue. Registration results from both tests are similar; the first is presented. The most recent trial of the actual treatment is presented. Handling of animals was in accordance to an Institutional Animal Care and Use Committee approved protocol. General anesthesia was induced, utilizing medetomidine (0.5 mg/kg, intramuscular) and 2% isoflurane was used to maintain general anesthesia. Axial and coronal planning images were acquired using a 1.5 T MR scanner (EXCITE HD GE Medical Systems, WI) with a GE 8 4-channel, receive-only phased-array body coil (MRI Devices Corp, Gainesville, FL, USA).

Pre-operative axial imaging data of the canine prostate and the neighboring anatomy was used to create a FEM mesh (Fig. 12). The pre-operative imaging data was taken at a resolution of $256 \times 256 \times 36$ voxels with a field of view of $240 \text{ mm} \times 240 \text{ mm}$ and an out of plane thickness of 3 mm with slice spacing of 1.5 mm. The FEM mesh consisted of 23,303 dof and 20,064 hexahedral elements. A pipeline of commercial software was used to generate the mesh. AMIRA, <http://www.amiravis.com>, was used to perform a semi-automatic segmentation of the prostate

and create a faceted surface representation of the prostate. CUBIT, cubit.sandia.gov, was used to convert the faceted surface into a hexahedral mesh and apply boundary conditions. Zero flux ($g = 0$) boundary conditions are used; heating is not expected to reach the outer boundary of the prostate. Intra-operative axial imaging data are shown in Fig. 12. The intra-operative imaging data was taken at a resolution of $256 \times 256 \times 12$ voxels with a field of view of $240 \text{ mm} \times 240 \text{ mm}$ and an out of plane spacing of 3 mm. Comparison of the pre-operative image and the intra-operative image reveals that the prostate has moved a significant amount which poses a significant registration problem. Current capabilities permit rigid registration only. The rigid registration is based on detecting the outer surface of the canine, not the prostate within. The rigid registration code⁴ found the out-of-plane position but failed in-plane due to the prostate movement. Consequently, the mesh was registered manually using AVS for interactive visualization. Furthermore, stringent time constraints on the experiment do not permit the use of the full resolution volume data sets for the rigid registration. The volume data sets must be subsampled to facilitate a reasonable execution time of the serial rigid registration code. However, for rigid registration, subsampling the imaging data by a factor of four in-plane does not lose much information on the outer boundary of the anatomy of the canine and significantly speeds up the registration from approximately 25 min to roughly 2 min with about 1 mm difference in the final result.

A planning template was placed on the perineum of the dogs. Coronal images, field of view $20 \text{ cm} \times 20 \text{ cm}$,

⁴ITK: <http://www.itk.org>.

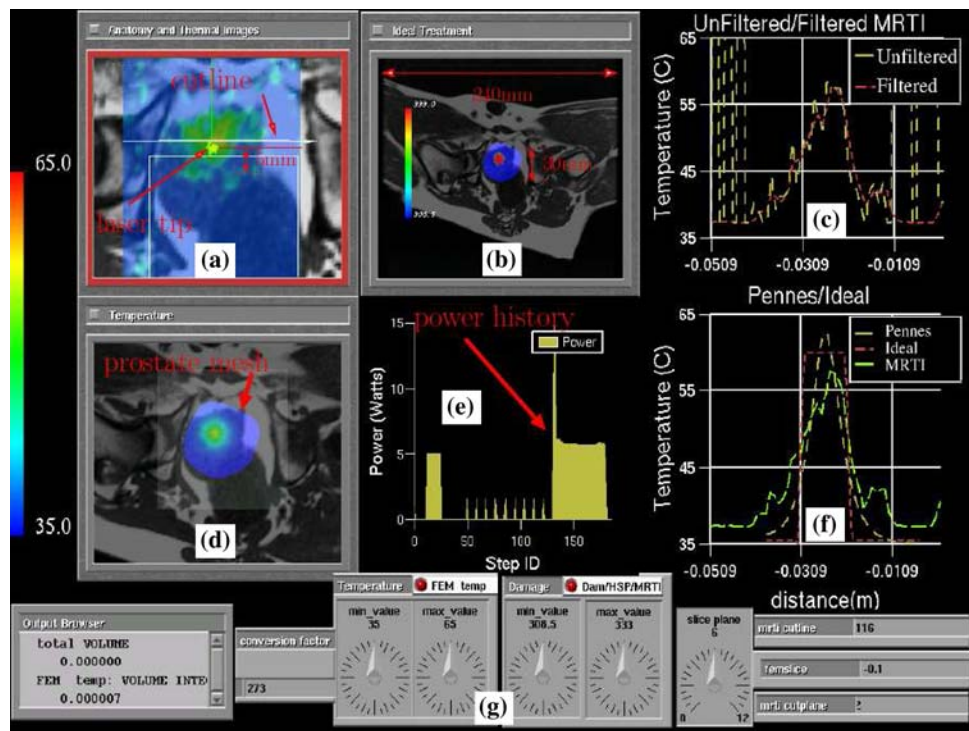


FIGURE 13. Depicted is the real-time visualization provided during the laser treatment. The color scale shown at the left is from 35–65 °C. (a) The laser tip may be seen within the anatomical images and the cutline through the thermal images use in (c) is shown. (b) The treatment goal was to create the 1.2 cm diameter lesion marked in red; for comparison, the diameter of the prostate is 30 mm and the field of view is 240 × 240 mm. (c) Temperature vs. distance along a cutline through the thermal images is shown. The noise in the unfiltered data, shown in yellow, has been removed in the filtered thermal signal, shown in red. The distance along the cutline is given in meters and the coordinates are with respect to the DICOM images from the scanner. (d) A contour plot of the FEM temperature prediction is shown on the prostate mesh. (e) The power history as a function of time is shown for the duration of the treatment. The history illustrates the four stages of the control system. The calibration pulse is followed by a series of keep-alive pulses of the laser. The keep-alive pulses prevent the laser from going into standby during the calibration and optimal control computations. The optimal power profile computed to create the lesion concludes the treatment. (f) Temperature vs. distance along a cutline through the FEM prostate mesh is shown. For comparison, the cutline is the same as used in (c). The treatment protocol was to create a lesion through a temperature history of 60 °C within the 1.2 cm diameter ball shown in (b). This is illustrated through the red cutline. The yellow cutline is the Pennes model prediction of the temperature along the cutline. Finally the green cutline is of the thermal images projected onto the finite element mesh. The temperature measured in the thermal images was controlled by supercomputers over a distance of 150 miles and shows within a 5 °C agreement at the peak. (g) A few of numerous controls needed to precisely display the remote visualization of the treatment are shown. The controls shown may be used to adjust the color scale of the temperature and the slice plane of the thermal/anatomical images as well as the position of the contour slice of the FEM mesh.

slice thickness 1.5 mm, covering the template and the prostate were input into the template planning program of the Visualase. Three “ducials on the template were identi“ed allowing trajectories to be projected through the prostate volume. A stainless steel stylet was used for inserting the laser catheters, 400 micron core diameter silica “ber in a water-cooled di used tip catheter, to reach the prostate at a depth of approximately 6 cm. The “nal catheter location was estimated using another set of coronal images. The estimated location of the laser, in DICOM coordinates, was given to HP3d to begin the treatment. A late time treatment image of the control system applied to in vivo laser treatment of canine prostate is shown irFig. 13; the entire duration of the treatment was 18 min. As shown, Fig.13b, the goal was to demonstrate that the computational model

could control the bioheat transfer and heat a region of 1.2 cm in diameter to 60°C. As illustrated in the power history, Fig. 13e, the treatment protocol is divided into the four stages described in Fig4. During the “rst stage, the biological domain is pulse heated and MRTI thermal image data are acquired for the heating as well as the cooling. The following parameters were used: “eld of view = 24 cm, slice thickness= 4 mm, echo train length = 8, acquisition matrix = 256 × 128, repetition time = 545 ms, echo time= 15 ms and “ip angle = 60° and “ve slices were collected per time point. The model calibration is divided into a 180 s data acquisition phase followed by a 180 s computation phase. For their vivo experiments, the canine prostate tissue was heated with a 5 W pulse for 90 s; a 180 s imaging period is used capture the cooling of the

in vivo prostate tissue as well as the heating. The acquired imaging data are used for model calibration computations. The treatment-day thermal imaging data generated for the canine exhibits significant decrease in the signal-to-noise ratio compared to those of the phantom experiments; the positioning grid separated the top coil from the surface of the canine's skin and separated it even further from its complimentary coil underneath the animal. The use of a calibration pulse of less than 5 W applied power on a previous experiment resulted in minimal heating, suggesting that the deposited energy was quickly dissipated by blood perfusion and any heating seen was masked by the noise range of the thermal imaging. The treatment imaging data shown in Fig. 13c uses a space-time filter; in addition to a spatial median-derivative filtering pipeline, if the thermal data at a pixel changes by more than 11°C it is considered noise and filtered. The 1°C threshold is based on intuition from previous in vivo experiments. The second stage of the treatment accounts for the time span of the calculations that use the imaging data for model calibration. The third stage accounts for the time delay to recompute the patient-specific optimal temperature/damage/HSP heating protocol. The laser is prevented from entering a standby mode through a series of keep-alive pulses

during the computational phase. In the fourth and final stage, the optimal laser control parameters are applied to the biological domain. For this particular trial, the perfusion, thermal conductivity, absorption coefficient, and laser position were calibrated. Results of the real-time calibration computations moved the laser tip 5 mm from the initial estimated position. Cutlines illustrating the temperature as function of distance are taken through the thermal imaging data and finite element predictions for comparison. The location of the cutline is illustrated in the Anatomy and Thermal Images (Fig. 13a). The bottom right cutline, Fig. 13f, shows good agreement between the desired treatment plan, the computation prediction, and the experimentally measured MRTI temperature field. A drift in the temperature imaging on the order of -1.7°C was seen over the course of the 18 min treatment time; implying that the thermal image value may be slightly higher than shown. However, the drift at the beginning of the treatment was insignificant during the 90 s thermal image acquisition for the calibration.

The prostate was excised immediately after exsanguination. It was fixed in formalin for 4 days, and then sectioned with guidance of the MR slice locations; this provided pathology slices which were congruent to the post-treatment MR images for comparison (Fig. 14).




FIGURE 14. (a) The damage, outlined by green arrowheads, is noted on a post-dynamic T1 image. This image corresponds to the central slice, $112 \pm 11 \text{ mm}^2$, of the thermal imaging, slice 3 of 5. The laser fiber is illustrated by the red arrow. Adjoining slices are also shown at (b) slice 4 of 5, $88 \pm 4 \text{ mm}^2$, and (c) slice 5 of 5, $42 \pm 4 \text{ mm}^2$. The lower row shows the damage on the corresponding sectioned pathology slices—(d) $110 \pm 5 \text{ mm}^2$, corresponds to (a), (e) $97 \pm 4 \text{ mm}^2$, corresponds to (b), and (f) $46 \pm 4 \text{ mm}^2$, corresponds to (c). A distance scale in centimeters is provided in (g).

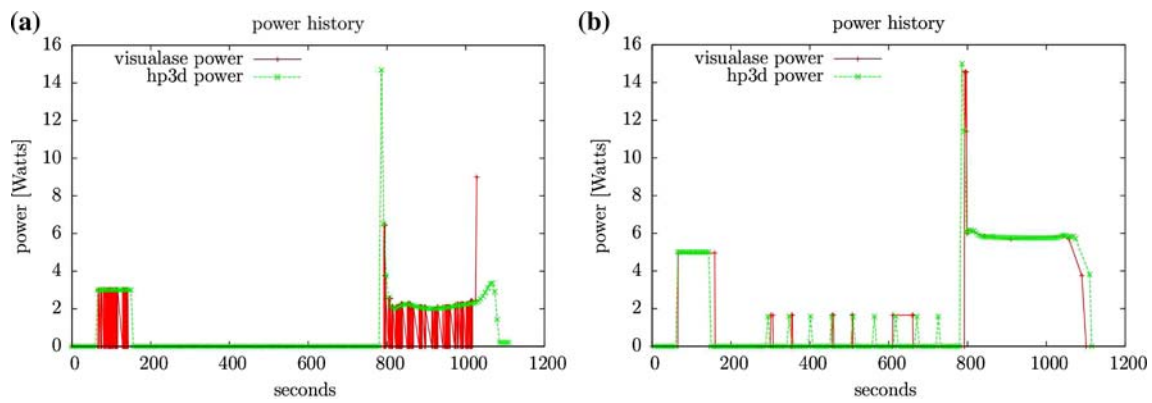


FIGURE 15. A comparison between the expected and actual power history as a function of time is shown for two canine experiments. The power history was extracted from the log files of the Visualase for comparison with the expected power profile from the FEM bioheat transfer code. The graphs show a time delay of approximately one image acquisition between the expected laser control and the actual laser control. (a) Copying the power control file directly to the mounted Visualase directory interferes with Visualase power log and instantaneously turns the laser off. (b) An implementation manipulating the file pointers was needed to keep the laser on.

The damage region was within the targeted 1.2 cm diameter treatment objective.

The actual power history recorded by the Visualase for the two experiments is shown in Fig. 15. The actual and predicted history are seen to be in very good agreement. However, the minor time delays seen result in a random single pulse not being physically manifested; but, post-treatment computations show that the delay seen in the power control is of little consequence. The 6-s 15 W spike missing from the Visualase log, Fig. 15a, has an insignificant effect on the model predictions. Figure 15a also shows that copying the power control file directly to the mounted Visualase directory interferes with Visualase power log and instantaneously turns the laser off. The problem is a result of the file system architecture. Overwriting by copying the file instantaneously deletes the file which instantaneously turns the power off. The fix is trivial, overwriting the file by moving does not delete the file and simply changes the file pointers (Fig. 15b). The sharp rise in power to 9 W at the end of the time history of the Visualase log file, Fig. 15a, is due to a manual override of the laser control during the earlier canine experiment.

CONCLUSIONS AND FUTURE DIRECTIONS

We believe that the results reported in this investigation support the proposition that robust simulation methodologies, based on calibrated nonlinear models of bioheat transfer in heterogeneous tissue, can be designed to interact with clinical thermal imaging modalities to provide near-real-time control of laser therapy for prostate cancer. The control system

employs what we refer to as a cyberinfrastructure: a computational grid connecting two remote arenas over a high bandwidth network. The imaging and laser treatment arena in Houston manages a MRTI device, a laser source, and relevant hardware and software units; the computational arena in Austin generates and solves a computational model of laser-driven bioheat transfer using parallel computing algorithms and visualization software and hardware to view the discrete system. New inverse analysis and optimization algorithms and software implementations of them are developed and used to control cell damage (cell viability), HSP-expression, and numerical accuracy. We employ a unified adjoint formulation for backward-in-time calculations for these control algorithms. The fact that numerical simulations and physical measurements interact in a feedback control paradigm in near real time qualifies this system as a DDDAS.

There are many factors that add to the complexity and difficulty in creating this particular DDDAS. The laboratory subjects used in our experimental trials were canines. Movement of the subject on the MRTI device before initiating therapy requires implementation of nonrigid registration algorithms to define the correct computational domain. Calibration of the heat transfer model is essential, and the presence of spatially heterogeneous thermal properties can lead to inverse analysis problems of enormous size. Numerical filtering algorithms are needed to filter noise inherent in the MRTI-generated thermal signatures. Accurate modeling and control of laser power and its correlation with temperature distributions is critical to the successful execution of the control.

Clearly, timing of the various events that evolve during laser therapy is very important and requires

access to large-scale parallel computing systems. In the trial runs described in this work, distributions of the computational work over 60...100 processors proved to be sufficient to reduce time-to-solution of the computer prediction to interval in which all of the control steps (calibration, power control, etc.) would be done without affecting the results. But for systems in which multiple calibrations of highly heterogeneous tissue are needed, 1000 or more CPUs may be needed. Algorithms to exploit the full power of forthcoming peta-scale computers are needed to solve the parabolic constrained optimization problems in a fraction of the time of the actual therapy. This will allow treatment protocol design that is able to calibrate and recompute optimal parameters instantaneously. Current work focuses on using an isotropic laser source to provide the heating and MRTI to provide the imaging in a thermal therapy for treating cancer. From a computational point of view, changing the thermal source or imaging modality merely amounts to changing the source term in the governing PDE or adjoint problem, respectively. This technology has the potential to be extended to many areas of thermal treatment, including RF, microwave, ultrasound, and even cryotherapy applicators. Modeling of the current laser source may also be improved by adding the ability for real-time Monte Carlo simulations to predict the laser fluence distribution.

Several objective functions have been presented in this work that mathematically characterize a desired treatment plan. It may be argued that different objective functions may be used to better represent the desired goal. For example, current damage functions account for only the damage accumulated during heating. Other objective functions may be designed to account for cooling as well as heating. The damage and HSP treatment models also need to be validated in vivo. Furthermore, different metrics may be explored to provide a better characterization of the distance between the computed and the ideal solutions. Finally, the objective function evaluation and gradient computation have been the workhorse in this project. Adding the functionality for the Hessian computation may enhance the speed and efficiency of the optimization steps for more complex bioheat models involving larger parameter sets and could add the ability to provide the framework for adding uncertainty quantification for the calibrated model parameters. Given the increasing ability of diagnostic radiology to detect diseases earlier and earlier, the impact of minimally invasive approaches to surgery, such as thermal therapy, will play a large role in safely and efficiently treating cancer with minimal impact on the patient. We hope that this work makes a small step toward the time at which computer modeling and simulation interacting

with medical technologies can dramatically improve cancer therapies and enhance and prolong the life of cancer patients.

ACKNOWLEDGMENTS

The research in this paper was supported in part by the National Science Foundation under Grants CNS-0540033, IIS-0325550, and NIH Contracts P20RR0206475, GM074258. During the course of this work we benefited from advice and comments of many colleagues, we mention in particular, C. Bajaj, J. C. Browne, I. Babuška, J. Bass, L. Bidaut, L. Demkowicz, Y. Feng, S. Goswami, A. Hawkins, S. Khoshnevis, B. Kwon, and S. Prudhomme. The authors also acknowledge the important support of DDDAS research by Dr. Frederica Darema of NSF. The authors would like to thank Drs. Ashok Gowda and Roger McNichols from BioTex Inc. for providing the Visualase[®] System and altering the software so that it may be remotely controlled from Austin, TX.

REFERENCES

- ¹A.V.S. Inc. Avs user's guide, National Congress on Computational Mechanics. San Francisco, CA, Conference Presentation, May 1992.
- ²Arbogast, T., and J. Bona. Methods of Applied Mathematics. University of Texas at Austin, Department of Mathematics, Applied Mathematics Course Notes, 2004.
- ³Atkins, D. E., K. K. Droegemeier, S. I. Feldman, H. Garcia-Molina, M. L. Klein, D. G. Messerschmitt, P. Messina, J. P. Ostriker, and M. H. Wright. Revolutionizing science and engineering through cyberinfrastructure: report of the national science foundation blue-ribbon advisory panel on cyberinfrastructure, Technical report, National Science Foundation, 2003.
- ⁴Balay, S., W. D. Gropp, L. C. McInnes, and B. F. Smith. Petsc users manual, Technical report. ANL-95/11,,Revision 2.1.5, Argonne National Laboratory, 2003.
- ⁵Benson, S. J., L. C. McInnes, J. Moré and J. Sarich. TAO user manual (revision 1.8), Technical report. ANL/MCS-TM-242, Mathematics and Computer Science Division, Argonne National Laboratory, 2005. <http://www.mcs.anl.gov/tao>.
- ⁶Darema, F. DDDAS: dynamic data driven applications systems. Website <http://www.nsf.gov/cise/cns/dddas>
- ⁷Demkowicz, L. Computing with hp-Adaptive Finite Elements: One and Two Dimensional Elliptic and Maxwell Problems, Vol. 1. Chapman & Hall/CRC Applied Sciences, 2007.
- ⁸Demkowicz, L., W. Rachowicz, and P. Devloo. A fully automatic hp-adaptivity. J. Sci. Comput. 17(1...3):127...155, 2002.
- ⁹Diller, K. R., J. T. Oden, C. Bajaj, J. C. Browne, J. Hazle, I. Babuska, J. Bass, L. Bidaut, L. Demkowicz, A. Elliott,

- Y. Feng, D. Fuentes, S. Goswami, A. Hawkins, S. Khoshnevis, B. Kwon, S. Prudhomme, and R. J. Stafford. Ch. 9: Computational Infrastructure for the Real-Time Patient-Specific Treatment of Cancer. In: *Advances in Numerical Heat Transfer: Numerical Implementation of Bioheat Models and Equations*, Vol. 3. Boca Raton, FL: Taylor & Francis Group, 2008.
- ¹⁰Diller, K. R., J. W. Valvano, and J. A. Pearce. Bioheat transfer. In: *The CRC Handbook of Mechanical Engineering*, 2nd ed., edited by F. Kreith and Y. Goswami. Boca Raton, FL: CRC Press, 2005, pp. 4-278...4-357.
- ¹¹Feng, Y., D. Fuentes, A. Hawkins, J. Bass, and M. N. Rylander. Model-based optimization and real-time control for laser treatment of heterogeneous soft tissue. *SIAM J. Sci. Comput.*, 2008, in review.
- ¹²Feng, Y., J. T. Oden, and M. Rylander. A statistical thermodynamics based cell damage models and its validation in vitro. *J. Biomech. Eng.*130(041016):1...10, 2008.
- ¹³Feng, Y., M. N. Rylander, J. Bass, J. T. Oden, and K. Diller. Optimal design of laser surgery for cancer treatment through nanoparticle-mediated hyperthermia therapy. *NSTI-Nanotech*2005(1):39...42, 2005.
- ¹⁴Oden, J. T., K. R. Diller, C. Bajaj, J. C. Browne, J. Hazle, I. Babuska, J. Bass, L. Demkowicz, Y. Feng, D. Fuentes, S. Prudhomme, M. N. Rylander, R. J. Stafford, and Y. Zhang. Dynamic data-driven finite element models for laser treatment of prostate cancer. *Num Meth PDE* 23(4):904...922, 2007.
- ¹⁵Pennes, H. H. Analysis of tissue and arterial blood temperatures in the resting forearm. *J. Appl. Physiol.* 1:93...122, 1948.
- ¹⁶Rylander, M. Design of hyperthermia protocols for inducing cardiac protection and tumor destruction by controlling heat shock protein expression. PhD thesis, The University of Texas at Austin, 2005.
- ¹⁷Rylander, M. N., Y. Feng, K. R. Diller, and S. J. Aggarwal. HSP27, 60, and 70 expression kinetics and cell viability following hyperthermia in normal and cancerous prostate cells. *Biophys. J.*, 2005, in review.
- ¹⁸Showalter, R. *Monotone Operators in Banach Space and Nonlinear Partial Differential Equations*. Providence: American Mathematical Society, 1997.
- ¹⁹Stafford, R. J., R. E. Price, C. J. Diederich, M. Kangasniemi, L. E. Olsson, and J. D. Hazle. Interleaved echo-planar imaging for fast multiplanar magnetic resonance temperature imaging of ultrasound thermal ablation therapy. *J. Mag. Resonan. Imag*20(4):706...714, 2004.
- ²⁰Welch, A. J., and M. J. C. van Gemert. *Optical-Thermal Response of Laser-Irradiated Tissue*. New York: Plenum Press, 1995.