Nonlinear Model Predictive Thermal Dose Control of Thermal Therapies: Experimental Validation with Phantoms

Dhiraj Arora, Mikhail Skliar, Daniel Cooley, Adam Blankespoor, Jeff Moellmer & Robert Roemer
University of Utah, Salt Lake City, UT

Abstract— **A novel approach to control of thermal therapies is evaluated experimentally using a one-dimensional agar phantom. According to the proposed approach, the thermal dose delivered to the target is controlled directly without attempting temperature control. The control problem is constrained by the transducer saturation constraints, and the constraints on the maximum allowable temperature in the normal tissue surrounding the target. The proposed thermal dose controller has a cascade structure with a linear constrained model predictive temperature controller in the inner loop, which receives the reference commands from a nonlinear thermal dose controller in the main loop. The controller is designed to allow for near time-minimal control of the treatment. A single, fixed, focused ultrasound transducer is used as a heating modality. The ultrasound power deposition and the perfusion in the phantom are identified experimentally, and used in the developed controller. The ability of the controller to deliver the desired thermal dose to the specified tumor region without violating temperature constraints in the surrounding normal tissue was evaluated using a series of experiments. The results demonstrate that the proposed approach is effective at delivering the desired thermal dose in a near minimal treatment time without violating constraints on the maximum allowable temperature in normal tissue. The experiments show that the controller can effectively operate even with substantial plant-model mismatch.**

I. INTRODUCTION

High temperature therapies, such as High Intensity Focused Ultrasound (HIFU) therapy, and ultrasound surgeries involve the use of elevated temperatures for therapeutic applications. The potential of focused ultrasound in treating tumors in human subjects has been demonstrated by several recent feasibility studies [1], [2]. During such thermal therapies, magnetic resonance allows precise target definition, location of focal spots and noninvasive temperature measurements. A simulation study by Vanne *et al.*, [3] showed the potential of using noninvasive MR temperature measurements as a feedback for online control of desired temperature trajectories in a specific region.

The effectiveness of thermal therapy treatments depends on the cumulative effect of elevated temperatures over the treatment time. The thermal dose quantifies the relationship between treatment time, temperature and treatment efficacy. However, due to the difficulties associated with the development of nonlinear thermal dose controllers, in the past, all treatment control approaches concentrated on controlling temperatures in several spatial locations instead of thermal dose. Because of the limited degrees of actuation freedom with existing transducers and transducer arrays, it is generally impossible to maintain the desired temperature distribution inside the target. At the same time, the thermal dose control problem is well posed, but is more difficult to implement because of a strongly nonlinear relationship between temperatures and the thermal dose. The control problem is further complicated by the transducer saturation constraints, and constraints on the maximum allowable temperatures in the surrounding normal tissue.

In this paper, the thermal treatment control problem is formulated and experimentally verified as a problem of controlling thermal dose. A cylindrical agar phantom was used to experimentally verify the model predictive controller developed in our earlier works [4], [5]. The ultrasound power deposition pattern and the perfusion in the phantom were obtained during a pre-treatment parameter identification. The obtained parameters are used in the internal model of the MPC thermal dose controller. The controller was used in real time to deliver a desired dose to a particular region of the phantom, designated as the tumor. The constraint handling capability of the MPC was used to constrain the maximum temperatures in the regions surrounding the "tumor", considered as normal tissue.

The results demonstrate that direct control of thermal dose is possible during thermal therapies. Furthermore, the near time-optimal nature of the controller is also verified, as proposed in our earlier works [4]–[5]. The results show that normal tissue temperatures can be constrained while delivering a desired thermal dose to the tumor.

II. THERMAL DOSE CONTROLLER

The block diagram of the developed thermal dose controller is shown in Figure 1, and is briefly described here for the case of a spatially distributed target. Further details of the controller design can be found in [4]–[5]. A limited number of temperature measurements, T_{mea}, are available in different spatial locations inside the phantom P, and used by the Kalman filter E to estimate the complete temperature distribution T_f inside the target and the surrounding tissue. Block S selects the T_{90} – the 10^{th} percentile of the estimated temperatures [6], which in turn is converted into delivered thermal dose in block H according to the model

D. Arora, D. Cooley and A. Blankespoor are with the Department of Mechanical Engineering.

M. Skliar is with the Department of Chemical Engineering. Send correspondence to: mikhail.skliar@utah.edu.

J. Moellmer is with School of Computing.

R. Roemer is with the Department of Radiology and Mechanical Engineering.

This work was supported in part by grants NIH-1-ROI-CA87785, NCI-ROI CA 33922 and NSF-CATS 0117300.

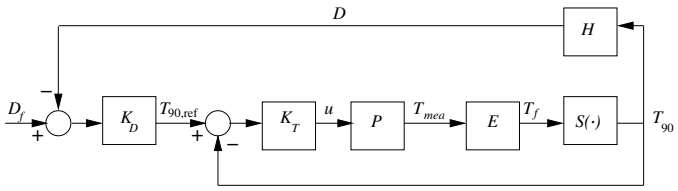

Fig. 1. Non-linear predictive thermal dose controller: K_D=Dose controller; K_T=MPC temperature controller; P=Phantom; E=Kalman estimator; S=T_{90} selector; H=Dose convertor.

proposed by Sapareto *et al.*, [7]:

$$D(t) = \text{CEM at } 43^{\circ}\text{C} = \int_0^t R^{[43-T_{90}(\tau)]} d\tau \qquad (1)$$

where $R = 0.5$ for $T_{90} > 43^{\circ}$C, $R = 0.25$ for $39^{\circ} < T_{90} < 43^{\circ}$C and $R = 0$ for $T_{90} < 39^{\circ}$C.

The constrained model predictive temperature controller K_T manipulates the intensity of the ultrasound transducer with saturation constraints to noninvasively heat the spatially distributed target. The predictive model used to characterize the temperature response in the agar phantom is described in section III. The objective function minimized by K_T is of the form

$$\begin{aligned} J(k) &= \sum_{j=1}^{p} w_y(j) \left[T_{90,ref}(k+j) - T_{90}(k+j)\right]^2 \\ &+ \sum_{j=1}^{m} w_u(j) \left[u(k+j-1)\right]^2 \end{aligned} \qquad (2)$$

where w_y is the penalty on the error between the reference and predicted T_{90}, and w_u is the penalty on the control effort. To prevent damage to the surrounding tissue, constraints on maximum allowable temperatures in the normal tissue are imposed (in experiments, the maximum allowable temperature of 4°C above baseline was imposed 14mm from the edge of the tumor). In our design, the control penalty is set to zero forcing the US intensity to its maximum limit u_{max}, except when constraints on the peak temperature in the normal tissue are about to be violated, as predicted by the thermal model. The model-based implementation of the temperature controller allows us to directly account for the thermal dose delivered during the cooling intervals when the transducer is switched off. To avoid infeasibility of the MPC problem, the constraints on the maximum allowable temperatures in the normal tissue are implemented as "soft" constraints. When satisfied, soft constraints have no effect of the objective function; when violated, a large penalty on the value of J is introduced.

The nonlinear thermal dose controller K_D in the main loop dynamically generates the reference temperature trajectories $T_{90,ref}$ for the predictive temperature controller. The reference temperature continuously generated by K_D is given by the following expression [4], [5]:

$$T_{90,ref}(t) = \frac{1}{\ln(1/R)} \ln \frac{\alpha(t_k)}{R^{43}}, \quad t \in [t_k, t_f] \qquad (3)$$

The subscripts k and f correspond to the current and the final treatment time, and the parameter α is calculated based on the difference between the delivered and the desired thermal dose, and the selected final treatment time:

$$\alpha(t_k) = \frac{(D_f - D(t_k))}{(t_f - t_k)} \qquad (4)$$

where $t_f = t_k + \Delta t$, with Δt being the tuning parameter of the thermal dose controller.

The thermal dose controller is designed to operate at the actuation or temperature constraints, which is theoretically required to achieve the minimum-time treatment control [4]. Selection of tuning parameter Δt governs the aggressiveness of the controller. A smaller value of Δt results in a more aggressive controller. The optimal value of Δt leads to the controller that operates at the power or temperature constraints. Further reduction of Δt may cause target overdose, which is usually acceptable from the clinical perspective. Larger values of Δt lead to a sluggish performance, resulting in treatments longer than optimal.

Theoretical study of the developed controller and the results of computer simulations with a one-dimensional model of a tumor were previously reported in [4], [5], [10]. These references provide further details on the properties of the developed controller, and its robustness and effectiveness at delivering the thermal dose.

III. PARAMETER IDENTIFICATION

The internal model used to predict the temperature response in a tissue was assumed to obey a 1-dimensional fin heat transfer equation

$$\rho C \frac{\partial T}{\partial t} = k \frac{\partial^2 T}{\partial x^2} - W(x) C [T - T_a] + Q(x) \qquad (5)$$

where C, k and ρ are the specific heat, thermal conductivity and density of the agar, respectively. In our experiments, it was assumed that C = 4186 J/(kg °C), k = 0.7 W/(m°C) and ρ = 1000 kg/m^3. The values of the parameters W(kg/m^3 sec) and Q(W/m^3), representing the convection heat transfer coefficient and the power deposition in the agar phantom, were obtained during "pre-treatment" parameter identification. These parameters were estimated by applying a step increase in the ultrasound power and recording the temperatures at discreet points along the length of the phantom. Assuming that conduction and convection effects are negligible immediately after the power is turned on ($t = 0+$), equation (5) simplifies to:

$$\rho C \frac{\partial T}{\partial t} = Q(x) \qquad (6)$$

The transient temperatures – starting from the baseline values (at $t = 0$) to the time corresponding to the end of the linear region, were used to compute the slope of the temperature increase at the each node. The slopes were used in equation (6) to obtain the power deposition at the corresponding nodes. Subsequently, a spline fit was used to obtain the complete power deposition pattern along the

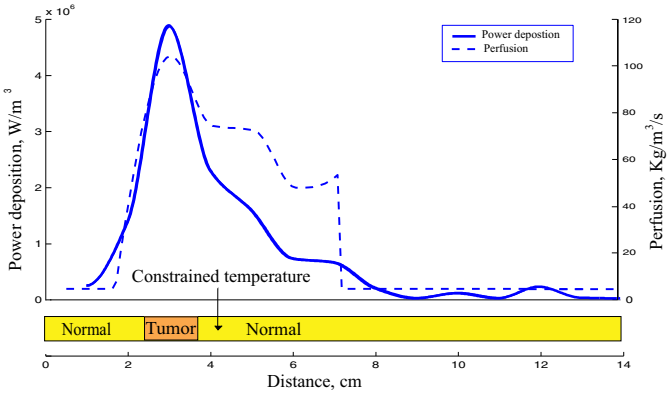

Fig. 2. Pre-treatment parameter identification: Power deposition pattern, $Q(x)$, and perfusion, $W(x)$.

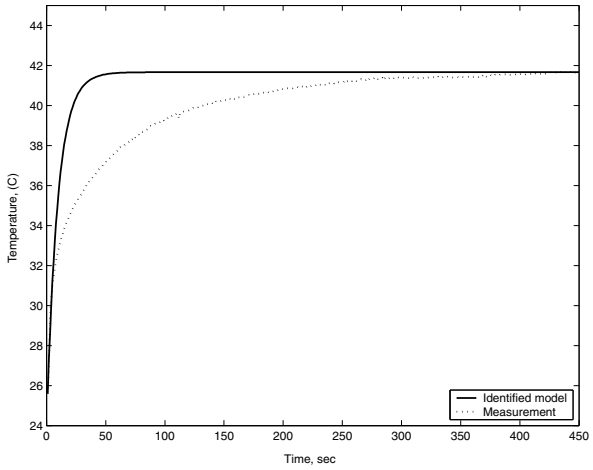

Fig. 3. Comparison of the model prediction with the measured temperature of the 3rd thermocouple

length of the agar phantom. Figure 2 depicts the identified power deposition of the focused ultrasound transducer operating at 1.5MHz and at the maximum allowable power (saturation constraints) of 16.3W. Further step test experiments showed that linear scaling can be used to obtain the power deposition curves for different applied powers.

To obtain the convection coefficient parameter, $W(x)$, the steady state temperatures from the step increase in power were used. At steady state, equation (5) reduces to:

$$k\frac{\partial^2 T}{\partial x^2} - W(x)C[T - T_a] + Q(x) = 0 \qquad (7)$$

The above equation and the thermocouple measurements were used to identify the values of a spatially varying convection coefficient, $W(x)$, with the results shown in Figure 2 as the perfusion. Figure 3 shows the comparison of the predicted temperature change in response to the step increase in the transducer power from zero to the maximum level of 16.3W obtained using the identified model with the measured values. The comparison for a single spatial point, corresponding to the location of the third thermocouple junction, clearly shows accurate steady state predictions, and considerable errors in the time constant of the response. The error in time constant is explained by the fact that the assumed fin heat transfer model (equation 5) is strictly one–dimensional, whereas, the agar phantom has a finite lateral dimension (1.27 cm in diameter). The conduction effects in the lateral dimension, which can be potentially significant for transducers with sharp focal zone, contribute significantly to the plant-model mismatch.

The 1-dimensional fin heat transfer model (equation 5) used to characterize the temperature response of the phantom is akin to the Pennes' bioheat transfer equation (BHTE) [8] except that the blood perfusion term in the BHTE is replaced by the convection term in the fin equation. In our experimental setup, the heat exchange with the water surrounding the phantom simulates the cooling effect of the blood perfusion in the tissue. Additional details on identification of the blood perfusion and the power deposition in a tissue can be found in [9].

IV. EXPERIMENTAL SETUP

The parameter identification and control experiments were performed on a 1-dimensional cylindrical agar phantom, 1.27cm in diameter and 15.5cm long. Condensed milk was added to the agar to achieve the ultrasound absorption similar to that of the human tissue. The transducer was placed in a bath of degassed and deionized water. The ultrasound and the thermometry system are shown in Figure 4.

The ultrasound field was generated by a single, spherically focused, air backed transducer resonating at a frequency of 1.5MHz. The radius of curvature and the diameter of the transducer were both 10 cm. An ultrasound positioning system was used to position the focal zone on the third thermocouple. Both the transducer and the positioning system were built in-house.

A function generator (Stanford Research System, Sunnyvale, CA, model DS345) and a radio frequency (RF) amplifier (ENI Inc, Rochester, NY, model A150) were used to generate the RF signal to drive the transducer. The electrical impedance of the transducer was matched to the output impedance of the amplifier by an external LC matching amplifier. The forward and backward reflected powers were measured using power meters (Hewlett-Packard, model 435A/B) and a dual directional coupler (Werlatone, model C625). The temperature in the phantom was monitored using two in-house built and calibrated seven-sensor manganin-constantan thermocouple probes. The first thermocouple was positioned 0.5 cm into the phantom and the remaining junctions were located one centimeter apart over the length of the phantom, Figure 5. An in house temperature acquisition system consisting of A/D cards and signal conditioning (block DAQ in figure 5) was used to acquire and monitor temperature measurements in 14 different spatial locations inside the agar.

The tumor was assumed to be located 22mm into the

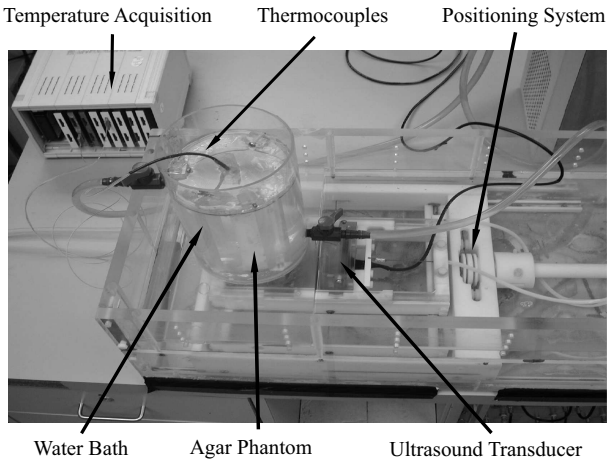

Fig. 4. Ultrasound positioning and temperature acquisition system

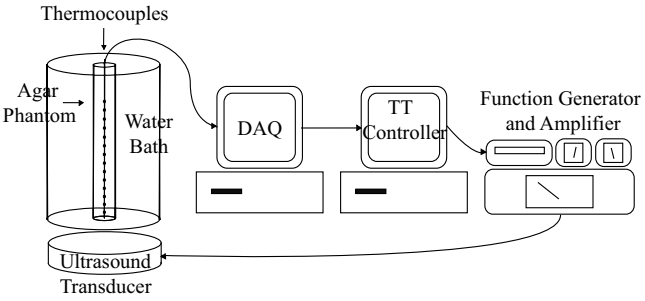

Fig. 5. Feedback control schematic for thermal therapy of a 1-D phantom

phantom, Figure 2, and is 10mm in length. Both constrained and unconstrained cases were considered. In the constrained case, the controller was required to maintain the temperature at the spatial position 44mm into the phantom (14mm from the edge of the tumor) below 41°C (or below 4°C above the baseline phantom temperature), Figure 2. The controller performance was adjusted using several tuning parameters, of which Δt is the most important.

During the experiments, the temperature measurements from the DAQ computer were transferred in real time to a dedicated control computer, shown as the TT controller. The computed control signal (transducer's power), measured and estimated temperatures, and the delivered thermal dose were monitored and stored in the TT control computer (Figure 5). The calculated control signal was used to drive the function generator using GPIB interface (National Instruments, model PCI-GPIB).

V. RESULTS

All the control experiments were carried out with an objective to deliver a dose of 240 cumulative equivalent minutes (CEM) of thermal dose to the selected "tumor" region of the phantom. Guided by the results from the simulation studies [4], [5], different values of the tuning

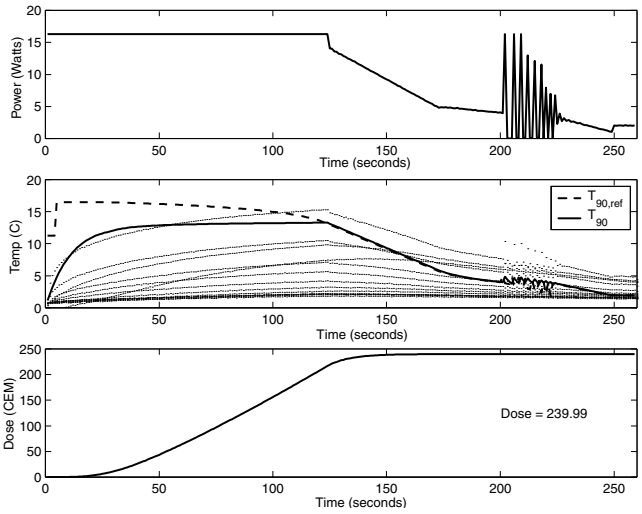

Fig. 6. Case of unconstrained temperature. $m = p = 1$ and $\Delta t = 10$. (a) Control input. (b) Increase in temperatures. Measured values are shown in grey. (c) Thermal dose.

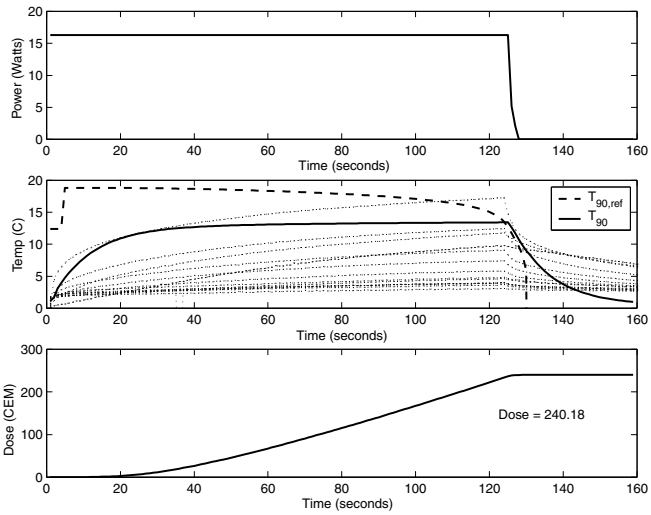

Fig. 7. Case of unconstrained temperature. $m = p = 1$ and $\Delta t = 2$. (a) Control input. (b) Increase in temperatures. Measured values are shown in grey. (c) Thermal dose.

parameter Δt were tested, and their effect on controller performance was analyzed. Cases with both constrained and unconstrained normal tissue temperatures were tested.

A. Unconstrained normal tissue temperature

Figure 6 depicts the results when both the prediction and the control horizon were set equal to unity, ($m = p = 1$) and $\Delta t = 10$. There were no constraints on the normal tissue temperature but a constraint on maximum allowed ultrasound power, $u_{max} = 16.3$W was present. The control input, Figure 6(a), stays at the maximum allowed value for most of the treatment and gradually reduces as the delivered dose reaches the setpoint value. In this example, longer than the optimal value of Δt is used, which results in a

less aggressive tuning of K_D, and therefore a control input that is not always at its constrained value. Towards the end of the treatment an additional power increase was required to deliver the desired thermal dose. Less aggressive tuning ensures the precise delivery of the reference thermal dose, as depicted in Figure 6(c). The temperature measurements from the 14 thermocouples as well the estimated T_{90} is shown in Figure 6(b). Also shown is the reference $T_{90,ref}$ sent by the dose controller K_D to the temperature controller K_T. The results show that the K_T is unable to track the reference at the beginning of the treatment, which is the intentional feature of the designed controller that forces the transducer operation at its maximum allowed power. However, as the delivered dose approaches the desired final dose, K_T tracks the reference faithfully. The oscillatory behavior of the control input towards the end of the treatment ensures that an exact dose of 240 minutes is delivered, and is allowed because in this application there is no practical reason to penalize the rate of change of the control signal.

A number of experimental runs were conducted to find the optimal value of Δt. Such a value of Δt will result in a minimum treatment time with no overdose. Figure 7 shows the results with $\Delta t = 2$ and $m = p = 1$. The control input is essentially maintained at one of its constraints, resulting in near minimum time treatment. The controller is switched from its upper constraint value at $t = 125$s and is switched off at $t = 128$s, with the residual thermal dose delivered during cooling of the target. Further decrease in Δt, corresponding to an even more aggressive controller will lead to an overdose of the tumor. Time optimal results of Figure 7 show substantial reduction in the treatment time compared to the suboptimal case of Figure 6.

B. Constrained normal tissue temperature

In the next set of experiments, in order to prevent thermal damage to the surrounding normal tissue, a constraint on the maximum allowable temperature (4°C above baseline phantom temperature) was imposed in spatial location indicated in Figure 2, 14mm from the edge of the tumor. Figure 8 depicts the results for the constrained case where $m = p = 1$ and $\Delta t = 10$. As evident from the Figure 8(b) and (c), the constraint is met but it was necessary to substantially prolong the treatment compared to the unconstrained case of Figure 5 in order to deliver the desired thermal dose. Since the rate of change of the control signal is not penalized, the temperature constraint is met with active modulation of the transducer power, which does not present any practical problem because of no mechanical components in the actuation system. Note that active switching is needed to achieve time-optimal treatment. For the ideal case of no plant-model mismatch (shown by simulations in our earlier work), the controller can exactly predict the power level that will maintain the normal tissue at the constrained value. Figure 9 depicts the results for the constrained case with Δt set equal to 2, which corresponds to a more aggressive controller tuning. In this case, the reference temperature,

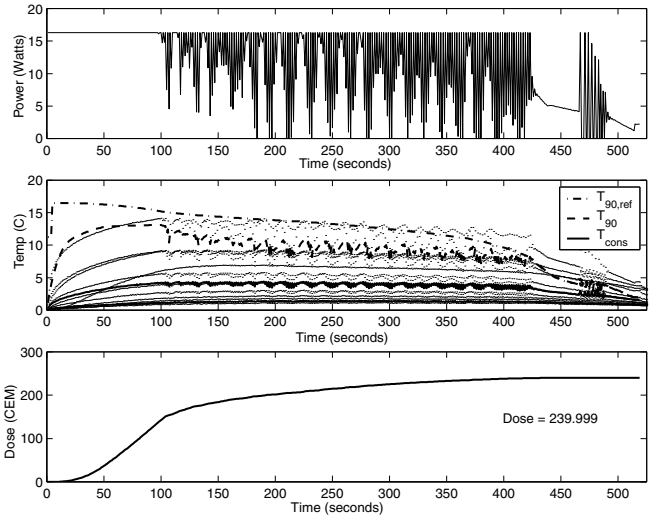

Fig. 8. Case of constrained temperature. $m = p = 1$ and $\Delta t = 10$. (a) Control input. (b) Increase in temperatures. (c) Thermal dose.

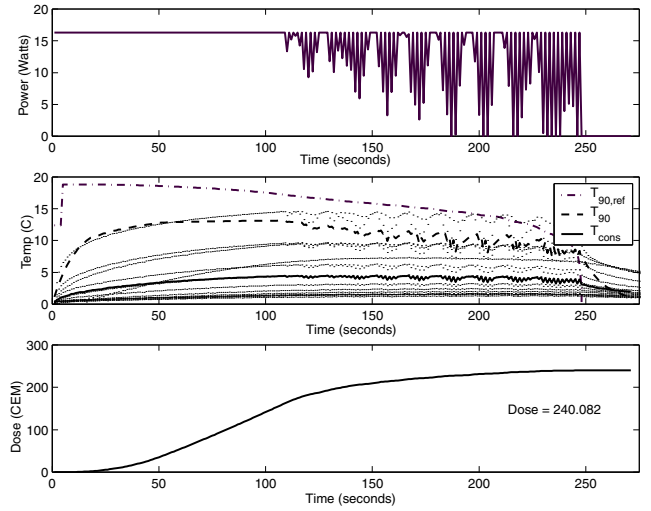

Fig. 9. Case of constrained temperature. $m = p = 1$ and $\Delta t = 2$. (a) Control input. (b) Increase in temperatures. (c) Thermal dose.

Figure 9(b), is higher than achievable with the given temperature and power constraints, forcing the controller to stay at either the maximum allowed power, or temperature constraint. Note that in this case, the controller operates closer to the maximum allowable constrained temperature, which significantly lowers the overall treatment time. The significant effect of operating closer to the temperature constraint on the overall treatment time is the consequence of a highly nonlinear relationship between temperatures and the thermal dose. The oscillatory nature of the control output is again due to the presence of the plant-model mismatch, and active control law that does not penalize aggressive control actions and allows for a high rate of change in control signal.

VI. DISCUSSIONS AND CONCLUSIONS

The control of thermal therapy has been traditionally formulated as a temperature control problem. We previously argued against such formulation because achieving the desired temperature distribution in the tumor with complex geometry, spatially varying blood perfusion and with the currently available power deposition actuators is an underactuated control problem. Instead, we proposed the control of thermal dose, which is a widely accepted clinical measure of the efficacy of thermal treatments. In contrast to the temperature control, the thermal dose control is a well posed problem. However, the design of the thermal dose controller is difficult because of the highly nonlinear relationship between temperature and thermal dose.

Our earlier works [10], [4] presented the first ever development and evaluation of different thermal dose controllers using simulations with a one-dimensional target. We demonstrated that in order to account for the thermal dose delivered during cooling, a model predictive controller is needed. Predictive control is especially important when multiple, interacting, high intensity pulses are used, yielding power superposition and high fluctuations of temperatures. To prevent normal tissue damage during high intensity thermal treatments, constraints on the maximum allowable temperature in the normal tissue must be imposed. Note that constraining maximum allowable thermal dose of the normal tissue will generally lead to noncausal controllers.

The developed controller solves the on-line optimization problem. Dynamic re-optimization of the treatment allows for robust delivery of the desired thermal dose despite modeling errors and process disturbances. Previous theoretical analysis and computer simulations show that the designed controller can be used to minimize the treatment time. According to the proposed approach, the aggressiveness of the minimum-time controller is balanced by imposing temperature constraints in normal tissue. Experimental results confirm previous analysis and agree well with our simulation studies.

The present work is the first experimental evaluation of a thermal dose controller, described in [4], [5]. The experiments are conducted with an agar phantom, which is modeled as a one-dimensional target. The power was deposited using a focused ultrasound transducer. The temperatures measured with 14 thermocouples, positioned on the centerline along the length of the phantom, are used in the feedback of the temperature controller. The available measurements and the identified thermal model are used in the Kalman filter to estimate the temperature profile inside the phantom. Based on the estimated profile the current thermal dose is calculated and compared with the desired final dose. The error is used as an input to the thermal dose controller.

The nonlinear predictive thermal dose controller has a secondary temperature and the primary thermal dose control loops. It was used during the experiments to deliver the desired thermal dose to the designated "tumor" region of the phantom. The thermal and power deposition models, required for the operation of the developed controller, are identified during pre-treatment stage based on the results of the ultrasound power step test. It is envisioned that in the clinical setting, the patient and site-specific thermal and power deposition models will be identified before the treatment using noninvasive magnetic resonance thermometry.

The experimental results show that the desired thermal dose is delivered to the tumor even in the presence of the a significant plant model mismatch, and without violating constraints on the normal tissue temperature and transducer saturation. Experiments show that the controller can be tuned to achieve the minimum time treatment – a critical factor for clinical acceptance of the thermal treatments.

It is expected that internal model predictions and the resulting controller performance can be further improved by using 3-D internal models in the controller. Among the challenges involved in development of 3-D thermal dose controllers are very large order systems and associated difficulties with real time implementation. It is known that the scanning and phased array transducers can significantly improve the precision of the TD delivery and reduce the treatment time. However, dynamic control of the scan trajectory and pattern re-phasing is a difficult and open problem currently under intensive investigation.

VII. ACKNOWLEDGEMENTS

The authors thank Trent Perry and Anthony Cummings, University of Utah, for their help with hardware development.

REFERENCES

[1] K. Hynynen, *et al.*, MR Imaging-guided Focused Ultrasound Surgery of Fibroadenomas in the Breast: A Feasibility Study, *Radiology*, vol. 219, 176–185, 2001.

[2] F. Wu, W.Z. Chen, J. Bai, Z.L. Wang, H.Zhu and Z.B. Wang, Pathological Changes in Human Malignant Carcinoma Treated with High-Intensity Focused Ultrasound, *Ultrasound Med. Biol.*, vol. 27, 1099–1106, 2001.

[3] A. Vanne & K. Hynynen, MRI Feedback Temperature Control for Focused Ultrasound Surgery, *Phys. Med. Biol.*, vol. 148, 31–43, 2003.

[4] D. Arora, M. Skliar and R. B. Roemer, Minimum Time Thermal Dose Control of Thermal Therapies, *IEEE Trans. Biomed. Eng.*, (submitted), manuscript available at `http://implicit.che.utah.edu/~skliar/BME2004.pdf`, 2003.

[5] D. Arora, M. Skliar and R. B. Roemer, "Nonlinear and Model Predictive Control of Thermal Dose in High Temperature Therapies", *in Proceedings of the American Control Conference*, Denver, CO, 1200–1205, 2003.

[6] K.A. Leopold *et al.*, Cumulative minutes with T_{90} greater than $Temp_{index}$ is predictive of response of hyperthermia and radiation, *Int. J Radiation Oncol. Biol. Phys*, vol. 25, 841–847, 1993.

[7] S.A. Sapareto & W.C Dewey, Thermal dose determination in Cancer Therapy, *Int. J. Oncology Biol. Phys.*, vol. 10, 787–800, 1984.

[8] H.H. Pennes, Analysis of Tissue and Arterial Blood Temperatures in Resting Human Forearm, *J. Appl. Physiol.*, vol. 1, 93–122, 1948.

[9] R.B. Roemer, A.M. Fletcher, and T.C. Cetas, Local SAR and Blood Perfusion Data from Temperature Measurements: Steady State and Transient Techniques Compared, *Int. J Radiation*, vol. 11, 1539–1550, 1985.

[10] D. Arora, M. Skliar and R. B. Roemer, Model-Predictive Control of Hyperthermia Treatments, *IEEE Trans. Biomed. Eng.*, vol. 49, 629–639, 2002.