

Mixed-integer model predictive control for large-area MR-HIFU hyperthermia in cancer therapy^{*}

D.A. Deenen^{*} E. Maljaars^{*} L. Sebeke^{**} B. de Jager^{*}
E. Heijman^{**,***} H. Gröll^{**} W.P.M.H. Heemels^{*}

^{*} Eindhoven University of Technology, Mechanical Engineering,
Control Systems Technology, Eindhoven, The Netherlands
(e-mail: d.a.deenen@tue.nl)

^{**} University Hospital of Cologne, Cologne, Germany

^{***} Philips Research Germany, Aachen, Germany

Abstract: In hyperthermia treatments, cancer tissue is heated to enhance the desired effects of radio- and chemotherapies. A powerful technology for noninvasive feedback-controlled hyperthermia is magnetic-resonance-guided high-intensity focused ultrasound (MR-HIFU), which enables fast and millimeter-accurate heating inside the body. Electronic beam steering allows for volumetric heating, but due to its limited steering range can only be used to treat small tumors. For the treatment of larger tumors, the transducer itself must be mechanically relocated as well. Due to system limitations, however, the admissible transducer positions must be restricted to a finite set that is chosen a priori. Moreover, non-negligible time is needed for transducer relocation, during which no heating is possible. In this paper, we present a mixed-integer model predictive controller that simultaneously optimizes over the power deposition by electronic beam steering – a continuous subproblem – as well as the mechanical transducer motions – a discrete subproblem. By incorporating model knowledge of the tissue’s thermal response and of the transducer carrier motion system into the predictive algorithm, the controller optimizes treatment temperature while respecting temperature and actuation constraints. The performance of the proposed feedback control setup is demonstrated by means of simulation.

Keywords: Model predictive control, mixed-integer programming, switched systems, hyperthermia, cancer treatment, large-area high-intensity focused ultrasound.

1. INTRODUCTION

In mild local hyperthermia treatments, cancer tissue is heated to temperatures around 42 °C for a duration of approximately 90 minutes. In the past decades, abundant clinical evidence has shown that mild hyperthermia acts on both cellular and tissue level, enhancing the cancer-killing effects of chemo- and radiotherapy, see Oei et al. (2015); Datta et al. (2016); Issels et al. (2018) and the references therein, for example. By locally heating only the tumor (and possibly some surrounding tissue), while avoiding temperature elevations in the remainder of the body, the tumor is selectively sensitized compared to the healthy tissue. Furthermore, using thermo-sensitive drug carriers, local hyperthermia enables the targeted release of anti-cancer drugs and improved tissue uptake specifically in the tumor area (Hijnen et al. (2014); Staruch et al. (2015)). Combined with the fact that hyperthermia itself introduces no additional toxicity, it allows for improved treatment outcome without increasing the undesirable side effects typically associated with cancer treatment,

^{*} This research has been made possible by the Dutch Cancer Society and the Netherlands Organisation for Scientific Research (NWO) as part of their joint Partnership Programme: “Technology for Oncology.” This project is partially financed by the PPP Allowance made available by Top Sector Life Sciences & Health. This work is also partially funded by the European Union via the IPaCT Project.

see Datta et al. (2015); Mallory et al. (2016). Therefore, hyperthermia is considered to be an effective and valuable addition to the existing cancer treatment modalities.

Controlling the target temperature in feedback can be done by means of magnetic-resonance-guided high-intensity focused ultrasound (MR-HIFU), where an MRI scanner is used for near-real-time thermometry, and an acoustic transducer generates ultrasound waves to locally deposit heat inside the body, see Maloney and Hwang (2015); Hynynen and Jones (2016). When an extracorporeal HIFU transducer is used, both the measurement and the actuation are completely noninvasive, which is highly desirable for patient comfort. Using a phased-array transducer, modulating the acoustic driver signals allows for rapidly moving the HIFU beam through the tumor. This is referred to as electronic beam steering, and enables powerful volumetric heating with millimeter-range accuracy.

In Deenen et al. (2018); Sebeke et al. (2019), model predictive control (MPC) has been successfully used to achieve safe and optimal heating by electronic beam steering for MR-HIFU hyperthermia. The key advantage of the developed MPC algorithms compared to other existing methods (e.g., the binary controller in Partanen et al. (2012) or the hybrid proportional-integral-derivative (PID) and bang-bang controller in Bing et al. (2015)) is their ability to explicitly take into account the tissue’s thermal dynamics

and the system's restrictive input constraints (such as the inability to actively remove heat from inside the body) when computing the desired heating input. Consequently, more homogeneous and stable target region temperatures can be achieved over a prolonged period of time.

Unfortunately, these MPC algorithms are currently only applicable to small-size tumors. That is, the accuracy of heating by HIFU using electronic beam steering comes at the cost of severely limited range, which hampers the treatment of larger tumors. To increase the treatable volume, the transducer itself can also be moved mechanically, see Tillander et al. (2016). Due to MR thermometry limitations, however, the admissible transducer positions must be confined to a discrete set. Moreover, mechanically relocating the transducer from one position to another requires a significant amount of time, compared to the sampling interval, during which no heating is possible. To allow for MPC to also enable the optimal treatment of larger tumors, the MPC designs that have previously been used to improve small-area MR-HIFU hyperthermia must be extended such that they can additionally take these effects into account.

Therefore, as the first key contribution of this paper, we develop a mixed-integer (MI-)MPC setup for large-area MR-HIFU hyperthermia. Our approach is specifically tailored to handle the switching of the transducer position and the thereby induced actuator downtime in a compact and systematic manner, leading to an efficient model (fitting in the mixed logical dynamical (MLD) system framework of Bemporad and Morari (1999)), that is directly suitable for the formulation of a constrained mixed-integer program (MIP). This MIP constitutes the core of our MI-MPC setup. The proposed controller solves on-line the constrained MIP, thereby simultaneously optimizing the local heating profile generated by electronic beam steering (while keeping the transducer stationary), which involves continuous control input variables, as well as the range-extending mechanical motion of the transducer to enable heating different parts of a large tumor, which is a discrete actuator allocation problem. When including the latter subproblem in the control algorithm, MPC again offers a major advantage over PID-based or binary controllers. That is, using MPC we are able to account for the significant transducer travel times when optimizing the transducer trajectory, which, due to the induced transducer inactivity, have a crucial influence on the achievable temperature distribution throughout the entire tumor. As the second main contribution of this paper, we demonstrate the performance of the novel control design by means of a realistic simulation study.

The remainder of the paper is organized as follows. First, the MR-HIFU hyperthermia setup and treatment are discussed in Section 2. Then, in Section 3 the state-space model describing the tissue's thermal dynamics is derived, where we pay specific attention to the manner in which the transducer switching is incorporated. In Section 4, we present the MI-MPC setup, for which we validate the effectiveness in achieving accurate temperature control for larger tumors in Section 5. Section 6 states the conclusions.

2. SYSTEM DESCRIPTION

In this section, we first discuss the MR-HIFU hyperthermia treatment setup, highlighting relevant aspects of the



Fig. 1. Philips 3T Achieva[®] and Profound Sonalleve[®] MR-HIFU therapy platform.

MR thermometry, and the heat delivery system consisting of the HIFU transducer and its carrier system. The latter enables the mechanical movement of the applicator. Then, we briefly describe the objectives of a hyperthermia treatment.

2.1 MR-HIFU hyperthermia therapy platform

Although the MPC setup developed in this paper is generally applicable to any system that exhibits actuator switches while the system dynamics remain unchanged, in this work it is designed for an MR-HIFU system as depicted in Fig. 1, consisting of the combination of a Philips 3T Achieva[®] MRI scanner used for near-real-time thermometry, and a Profound Sonalleve[®] HIFU platform. The latter is a dedicated trolley-tabletop in which an MR-compatible HIFU transducer and its carrier system are integrated, allowing for noninvasive heating of internal tissue volumes. The combined setup is already being used in clinics for the treatment of patients with uterine fibroids and for palliative pain management in patients with bone metastases.

2.2 MR thermometry

The thermal maps are obtained noninvasively using the temperature-dependent proton resonance frequency shift (PRFS) method, see Ishihara et al. (1995); Winter et al. (2016). This provides the temperature change relative to some baseline, by inspecting the difference of the current MR image with respect to a reference image. Typically, the reference is acquired before treatment when the treatment-induced temperature elevation is known to be zero.

Although this is one of the best available measurement technologies for image-guided hyperthermia feedback control, as argued in Quesson et al. (2000); Rieke and Butts Pauly (2008); Hynynen and Jones (2016), it is sensitive to modifications in the magnetic field. One source that can lead to significant distortion of the magnetic field is the transducer itself, even when it consists of non-ferromagnetic material. As a result, for accurate PRFS-based MR thermometry given a certain transducer position, a baseline image is required which was obtained with the transducer in the same position. Therefore, we constrain the admissible mechanical transducer positions to a predefined discrete set, such that we can obtain a reference image for each location before treatment, and use a lookup table to select the correct reference image for relative thermometry during treatment.




Fig. 2. Schematic of HIFU beam into focal plane in tumor area, with focal point by electronic beam steering.

2.3 HIFU transducer

The Sonalleve contains a phased-array HIFU transducer consisting of 256 acoustic elements, each of which is able to generate high-intensity ultrasound waves. By coordinated phase and amplitude shifts of the individual elements, referred to as electronic beam steering, a focal spot with variable position can be created at which the acoustic energy is highest, see Fig. 2. The power deposition profile is relatively large in the beam's axial direction, but narrow in radial direction. Consequently, the resulting volumetric temperature distribution is fairly uniform along the beam axis (sufficiently near the focal plane), meaning volumetric heating of the target tissue can be achieved by setting the focal plane halfway the target volume and controlling only the in-plane temperature distribution. Therefore, in this paper we consider temperature control in the two-dimensional focal plane only.

For a fixed transducer position, the treatable region within the range of electronic beam steering is referred to as a treatment cell. Using electronic beam steering, the focal spot can rapidly be scanned through the cell area, effectively allowing for temporally quasi-continuous heating of the tissue inside the cell. For lateral beam deflections larger than approximately 8 mm from the transducer axis, the focus quality significantly deteriorates. Consequently, a treatment cell cannot be larger than a 16 mm diameter circle.

The HIFU transducer is carried by a robot arm, which is also embedded in the MR-HIFU treatment table. Using the robot arm to mechanically displace the transducer allows for the heating of multiple cells, thereby increasing the treatable target size. Due to the aforementioned limitations of MR thermometry, the number of admissible treatment cell locations is finite, and they must be chosen before treatment. Moreover, once a transducer motion is initiated, it must be completed. In this paper, we desire to heat the region of interest (ROI) \mathcal{R} depicted in Fig. 3 using the $N_q = 4$ treatment cells also indicated in the figure. Here, the cell locations are chosen manually for simplicity. As future work, this may be optimized for improved control performance.

2.4 Hyperthermia treatment

The main goal in a hyperthermia treatment is to maintain a controlled and homogeneous temperature elevation in the ROI \mathcal{R} . The heat-induced tissue sensitization is triggered at temperatures around 41 °C, which is why the healthy tissue outside \mathcal{S} in Fig. 3 will be safeguarded against temperatures above 40 °C. Optimal treatment quality is achieved at 42 °C in the ROI, while overheating above 43 °C reduces the beneficial heat-induced effects and should therefore be avoided.

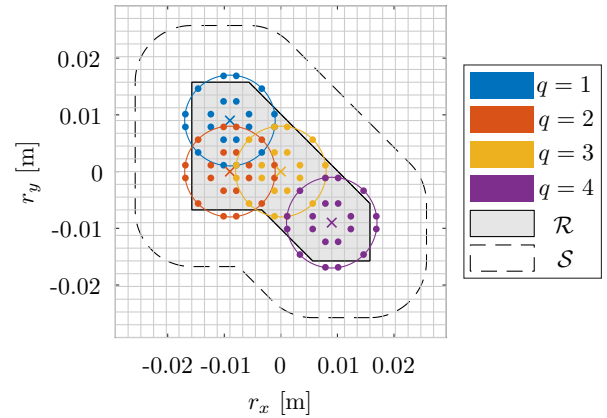


Fig. 3. The distribution of the sonication points (\bullet) of the four treatment cells, for which the center (\times) and electronic beam steering range (circles) are indicated, over a large ROI \mathcal{R} . The tissue outside \mathcal{S} must be safeguarded against overheating.

3. MIXED-INTEGER MODELING

This section discusses the state-space model describing the tissue's thermal dynamics, in which we incorporate the transducer switching in a compact and intuitive manner specifically designed to obtain a small-size MLD model that is directly suitable to incorporate in a MIP.

3.1 Transducer switching

In our modeling approach, we need only define the *operational* actuator modes

$$q \in \mathcal{Q} = \{1, \dots, N_q\}, \quad (1)$$

where N_q is the number of treatment cells. In Fig. 3, we have $N_q = 4$. Contrary to using a lifting approach (Subramanian et al. (2012)), which results in a constrained switched linear system (Philippe et al. (2016)), using our approach we need no additional “inactive” modes to model the actuator downtime during transducer switching. Instead, we will incorporate this behavior using systematically derived inequality constraints, resulting in a significantly more compact model. These constraints will be based on $\sigma_k \in \mathcal{Q}$, which denotes the *desired* (operational) actuator mode, i.e., the actuator mode which the system is in or towards which it is transitioning, at some time instant $k \in \mathbb{N}$. The time instant k connects to real time $t_k = kT_s$, where $T_s = 3.2$ s denotes the sample time that follows from the MR thermometry protocol. Note that σ_k effectively lumps together two “transducer states” (moving towards and being at a cell), which is instrumental in facilitating the reduced-complexity system description. Furthermore, to describe the cell-to-cell travel times using our approach, we need only conveniently specify the so-called setup time matrix, which for the example in Fig. 3 reads

$$S = \begin{bmatrix} 0 & 1 & 2 & 3 \\ 1 & 0 & 1 & 3 \\ 2 & 1 & 0 & 2 \\ 3 & 3 & 2 & 0 \end{bmatrix}, \quad (2)$$

of which the elements $s_{q\bar{q}}$ correspond to the travel time from cell q to \bar{q} , for any $q, \bar{q} \in \mathcal{Q}$. Since staying in a cell requires no movement and thus no travel time, we have $s_{qq} = 0$ for all $q \in \mathcal{Q}$. Moreover, $S = S^T$ in (2) since in the considered MR-HIFU system the travel times do not

depend on the relative direction of motion. Note, however, that our modeling approach can also be used if S is not symmetric.

3.2 Thermal state-space model

To model the tissue's thermal dynamics, we follow a procedure similar to Deenen et al. (2018). In summary, this entails first describing the underlying physical system using the Pennes bioheat equation (Pennes (1948)) and then spatially discretizing this partial differential equation on a two-dimensional square grid with $2.25 \times 2.25 \text{ mm}^2$ voxels (depicted in Fig. 3 by the grid lines) using the central difference scheme, which in Sebeke et al. (2019) was verified to be the method that best balances model simplicity with descriptive accuracy. Finally, the model is temporally discretized using the forward Euler method, which preserves model sparseness, while providing sufficient accuracy for the considered system.

The resulting discrete-time state-space model is given by

$$x_{k+1} = Ax_k + Bu_k, \quad (3a)$$

$$y_k = x_k + v_k, \quad (3b)$$

where the states $x_k \in \mathbb{R}^{n_x}$, $n_x = 36^2 = 1296$, represent the temperature elevations with respect to the baseline of the voxels in the focal plane. The matrix A has the form of a 2D discrete Laplacian matrix, capturing the effects of heat transfer by conduction and by blood perfusion. The voxels are chosen such that their centers coincide with the points measured by MR thermometry, resulting in full state measurements $y_k \in \mathbb{R}^{n_x}$ corrupted by MR measurement noise $v_k \in \mathbb{R}^{n_x}$, which can be well approximated by zero-mean Gaussian noise with a standard deviation of $0.4 \text{ }^\circ\text{C}$. Per treatment cell $q \in \mathcal{Q}$, we use $n_u^q = 20$ voxels at the centers of which we allow sonication by appropriate steering of the focal spot. These locations are referred to as sonication points, which are also shown in Fig. 3 by the markers '•'. The modal input $u_k^q \in \mathbb{R}_{\geq 0}^{n_u^q}$ describes the average sonication power applied at each of the sonication points in cell $q \in \mathcal{Q}$ over the course of one sampling interval. The modal inputs are collected in the input vector

$$u_k = \begin{bmatrix} u_k^1 \\ \vdots \\ u_k^{N_q} \end{bmatrix} \in \mathbb{R}_{\geq 0}^{n_u}, \quad (4)$$

where $n_u = \sum_{q \in \mathcal{Q}} n_u^q = 80$. Accordingly, the input matrix is given by

$$B = [B^1 \ \dots \ B^{N_q}] \in \mathbb{R}^{n_x \times n_u}, \quad (5)$$

where each submatrix $B^q \in \mathbb{R}^{n_x \times n_u^q}$ describes the system's thermal response to the heating power u_k^q at the sonication points in the corresponding cell $q \in \mathcal{Q}$.

3.3 Mixed-integer linear input constraints

First, for compatibility with the MLD modeling framework, let us formalize the nonnegativity of the heating inputs in (4) using the element-wise inequality

$$0_{n_u} \leq u_k, \quad \text{for all } k \in \mathbb{N}, \quad (6)$$

where 0_{n_u} denotes the zero vector of length n_u . Next, to describe that heating is only possible in the cell corresponding to the current transducer position, and impossible during transducer motion, we use inequality con-

straints on the input u_k that also involve logical variables. To this end, let us introduce the Boolean activation signal

$$\delta_k = \begin{bmatrix} \delta_k^1 \\ \vdots \\ \delta_k^{N_q} \end{bmatrix} \in \{0, 1\}^{N_q}, \quad (7a)$$

where

$$\delta_k^q = \begin{cases} 1, & \text{if } \sigma_k = q, \\ 0, & \text{otherwise,} \end{cases} \quad (7b)$$

for all $q \in \mathcal{Q}$, i.e., $\delta_k^q = 1$ if at time $k \in \mathbb{N}$ the transducer is at or traveling towards location $q \in \mathcal{Q}$ (indicated by $\sigma_k = q$), and $\delta_k^q = 0$ otherwise. This results in the identity

$$\sum_{q \in \mathcal{Q}} \delta_k^q = 1. \quad (8)$$

As previously mentioned, using our new modeling approach we do not need to introduce any additional auxiliary variables to explicitly describe the transducer inactivity during relocation. Instead, we infer whether the transducer is moving at time k using the information in S (2) in relation to the current and past activation signal values $\delta_{k-\tau}$ (which are directly related to $\sigma_{k-\tau}$), $\tau = 0, \dots, \bar{s}$, where $\bar{s} = \min\{k, \max_{q, \bar{q} \in \mathcal{Q}} s_{q\bar{q}}\}$ denotes the largest occurring setup time specified in S (upper bounded by k). The $N_q(\bar{s} + 1)$ input constraints that capture the transducer switching are then given by

$$I_{N_q n_u} u_k \leq \bar{u}_\Sigma \delta_k, \quad (9a)$$

$$I_{N_q n_u} u_k \leq \bar{u}_\Sigma S_\tau^\top \delta_{k-\tau}, \quad \tau = 1, \dots, \bar{s}, \quad (9b)$$

for all $k \in \mathbb{N}$, where $\bar{u}_\Sigma = 100 \text{ W}$ is the maximum allowable sonication power delivered during one sampling period, the matrices

$$S_\tau = \begin{bmatrix} s_{\tau,11} & \cdots & s_{\tau,1N_q} \\ \vdots & \ddots & \vdots \\ s_{\tau,N_q1} & \cdots & s_{\tau,N_qN_q} \end{bmatrix} \in \{0, 1\}^{N_q \times N_q}, \quad (10a)$$

follow uniquely from S in (2) according to

$$s_{\tau,q\bar{q}} = \begin{cases} 1, & \text{if } s_{q\bar{q}} < \tau, \\ 0, & \text{otherwise,} \end{cases} \quad (10b)$$

and $I_{N_q n_u}$ represents the block diagonal matrix

$$I_{N_q n_u} = \begin{bmatrix} 1_{n_u^1}^\top & & \\ & \ddots & \\ & & 1_{n_u^{N_q}}^\top \end{bmatrix} \in \{0, 1\}^{N_q \times n_u}, \quad (11)$$

such that $I_{N_q n_u} u_k$ is a vector of length N_q of which the q -th element $1_{n_u^q}^\top u_k^q$ describes the total input power in cell $q \in \mathcal{Q}$. Since each individual input element must be nonnegative by (6), constraining a cell's total input power in (9) to zero effectively disables all heating in that cell. To further clarify the working principle of the inequality constraints (9), let us consider the following example related to Fig. 3.

Example 1. Consider the input u_K^1 corresponding to cell $q = 1$ at some time $K \in \mathbb{N}_{\geq 3}$. Using S in (2), we compute the matrices S_τ , $\tau = 1, 2, 3$, as in (10) to be

$$S_1 = I_4, \quad S_2 = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad S_3 = \begin{bmatrix} 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}. \quad (12)$$

Then, by the inequality constraints (9), combined with the nonnegativity constraints (6), heating in cell 1 can occur only if it holds that

$$\delta_K = \delta_{K-1} = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad \delta_{K-2} \in \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} \right\}, \quad (13)$$

$$\delta_{K-3} \in \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} \right\}.$$

Note that (13) includes all sequences that allow for heating in cell 1 according to the travel times in S (2). In particular, (13) contains the cases that the transducer was:

- at or traveling towards cell 1 at all times $k = K - 3, \dots, K$, which combined with the largest travel time towards cell 1 being 3 samples, see (2), guarantees that the transducer is at cell 1 at time K at the latest;
- at or traveling towards cell 2 until time $K - 2$ at the latest, then traveling towards cell 1 at time $K - 1$, which by (2) requires 1 sample, and thus is at cell 1 at time K ;
- at or traveling towards cell 3 until time $K - 3$ at the latest, then traveling towards cell 1 at times $k = K - 2, K - 1$, which by (2) requires 2 samples, and thus is at cell 1 at time K .

The cases listed in Example 1 describe all activation signal sequences $\delta_{[K-3,K]} = (\delta_{K-3}, \dots, \delta_K)$ that are *admissible*, i.e., correspond to a physically realizable transducer path, and that allow for nonzero heating in cell 1 at time K . However, the inequalities (9) also allow for *spurious* sequences $\delta_{[K-3,K]}$ that do not describe a realizable trajectory, as can be seen from the combinations in (13) that do not correspond to one of the listed cases. The most obvious example is when no heating is required at some time $K \in \mathbb{N}_{\geq 3}$, i.e., $u_K = 0$, by which the values δ_k , $k = K - 3, \dots, K$, may differ at each time instant. Such sequences are undesirable, because they do not show, in an unambiguous manner, which cell the transducer should be at or traveling towards, and consequently cannot be used directly for mechanical transducer steering. Using our modeling approach, however, we need not introduce additional constraints to ensure admissibility of $\delta_{[0,K]}$, the derivation of which is not straightforward, but instead we use the procedure presented in the next subsection.

3.4 Admissibility assurance

First, note that (9) correctly constrains the control inputs in relation to *all* admissible transducer paths. Secondly, given any *feasible* pair $(\tilde{\delta}_{[0,K]}, \mathbf{u}_{[0,K]})$, i.e., where $\tilde{\delta}_{[0,K]}$ and $\mathbf{u}_{[0,K]} = (u_0, \dots, u_K)$ satisfy the inequalities (9) for all $k \in \mathbb{N}_{[0,K]}$, we can construct a (possibly different) sequence $\delta_{[0,K]}$ that is admissible and that renders $(\delta_{[0,K]}, \mathbf{u}_{[0,K]})$ feasible. In particular, one such (alternative) $\delta_{[0,K]}$ follows from a feasible pair $(\tilde{\delta}_{[0,K]}, \mathbf{u}_{[0,K]})$ according to

$$\delta_k = \begin{cases} \tilde{\delta}_k, & \text{if } u_k \neq 0, \\ \tilde{\delta}_{L_+(k)}, & \text{if } u_k = 0 \text{ and } L_+(k) \in \mathbb{N}_{[k+1,K]}, \\ \tilde{\delta}_{L_-(k)}, & \text{if } u_k = 0 \text{ and } L_+(k) \notin \mathbb{N}_{[k+1,K]}, \end{cases} \quad (14a)$$

for all $k \in \mathbb{N}_{[0,K]}$, where

$$L_+(k) = \inf\{l \in \mathbb{N}_{[k+1,K]} \mid u_l \neq 0\}, \quad (14b)$$

$$L_-(k) = \max\{l \in \mathbb{N}_{[0,k-1]} \mid (u_l \neq 0 \text{ or } l = 0)\}, \quad (14c)$$

which we refer to as the admissibility assurance. This is a computationally fast procedure, which effectively

constructs the admissible transducer path that enables the heating described by $\mathbf{u}_{[0,K]}$ while switching as little and as early as possible. Hence, after first solving an MI-MPC optimization problem that includes the inequality constraints (6) and (9) to find an optimizing feasible pair $(\tilde{\delta}_{[0,K]}, \mathbf{u}_{[0,K]})$, we can use (14) to swiftly obtain the admissible activation signal sequence corresponding to $\mathbf{u}_{[0,K]}$ to be used for the mechanical steering of the transducer.

4. MIXED-INTEGER MPC

In this section, we first state the controller and observer models. Then, we translate the MR-HIFU hyperthermia treatment objectives into a cost function and a set of constraints, which together with the admissibility assurance constitute our MI-MPC setup.

4.1 Controller and observer models

The MPC prediction model is given by

$$x_{i+1|k} = Ax_{i|k} + Bu_{i|k}, \quad (15)$$

where $x_{i|k} \in \mathbb{R}^{n_x}$ and $u_{i|k} \in \mathbb{R}^{n_u}$ denote the predicted states and inputs, respectively, at $i \in \mathbb{N}$ time steps ahead of the prediction sequence's starting time $k \in \mathbb{N}$. Although model uncertainty in (3) is typically inevitable in practice, for the scope of this work a deterministic approach suffices. For an indication of the influence of model error on control performance for small-size tumors, see Sebeke et al. (2019). As future work, the algorithm developed here may be adapted to be more robust, for example using an offset-free implementation such as proposed in Deenen et al. (2018) for single-cell MR-HIFU hyperthermia.

To obtain improved temperature estimates with respect to the noise-corrupted MR thermometry readings, we design a Luenberger-type observer given by

$$\hat{x}_k = A\hat{x}_{k-1} + Bu_{k-1} + L(y_k - \hat{y}_k^-), \quad (16a)$$

where

$$\hat{y}_k^- = A\hat{x}_{k-1} + Bu_{k-1} \quad (16b)$$

denotes the model-based measurement estimate at time k before applying the correction by output injection to obtain \hat{x}_k . The observer gain matrix is given by $L = 0.25I_{n_x}$, which yields a stable estimator, and has been tuned in simulation to achieve desirable convergence properties.

4.2 Optimization problem

The temperature objectives are schematically depicted in Fig. 4 in cross-section perspective. We use $\Omega \subset \mathbb{R}^2$ to denote the patient domain in the focal plane. $T_r : \mathcal{R} \rightarrow \mathbb{R}$ represents the optimal ROI treatment temperature. $\bar{T} : \Omega \rightarrow \mathbb{R}$ is the location-dependent desired upper temperature bound, which describes the fact that heating the ROI above 43 °C reduces the beneficial hyperthermia-induced effects, and that the region $\Omega \setminus \mathcal{S}$ should not be heated above 40 °C to avoid sensitization of healthy tissue. To translate these objectives to the state of (3), let us denote the states corresponding to points within \mathcal{R} by $z_k = Hx_k \in \mathbb{R}^{n_z}$, where $n_z < n_x$ and $H \in \{0, 1\}^{n_z \times n_x}$ is a matrix with exactly one 1 in each row (and at most one 1 per column). Also, let $z_r \in \mathbb{R}^{n_z}$ and $\bar{z} \in \mathbb{R}^{n_x}$ denote the voxel-wise temperature elevation reference and upper bound corresponding to the values




Fig. 4. Schematic cross-section of the temperature objectives corresponding to \mathcal{R} and \mathcal{S} . The maximum violation $\bar{\epsilon}$ is shown for some overheated temperature distribution T_1 such that $T_1 \leq \bar{T} + \bar{\epsilon}$.

of T_r and \bar{T} , respectively. The maximum violation of the upper temperature bound can then be described by the slack variable $\bar{\epsilon}_k = \bar{\epsilon}(x_k) = \|\max\{x_k - \bar{x}, 0_{n_x}\}\|_\infty \in \mathbb{R}_{\geq 0}$, where 0_{n_x} denotes a zero-vector of length n_x and the maximum operator is used element-wise.

The optimization problem is then given by

$$\min_{\delta_k, \mathbf{u}_k} \sum_{i=0}^N (z_{i|k} - z_r)^\top Q (z_{i|k} - z_r) + f_\epsilon \bar{\epsilon}_{i|k}, \quad (17a)$$

subject to

$$x_{i+1|k} = Ax_{i|k} + Bu_{i|k}, \quad \forall i \in \mathbb{N}_{[0, N-1]}, \quad (17b)$$

$$x_{0|k} = \hat{x}_k, \quad (17c)$$

$$x_{i|k} \leq \bar{x} + 1_{n_x} \bar{\epsilon}_{i|k}, \quad \forall i \in \mathbb{N}_{[0, N]}, \quad (17d)$$

$$0 \leq \bar{\epsilon}_{i|k}, \quad \forall i \in \mathbb{N}_{[0, N]}, \quad (17e)$$

$$0_{n_u} \leq u_{i|k} \leq 1_{n_u} \bar{u}, \quad \forall i \in \mathbb{N}_{[0, N-1]}, \quad (17f)$$

$$I_{N_q n_u} u_{i|k} \leq \bar{u}_\Sigma \delta_{i|k}, \quad \forall i \in \mathbb{N}_{[0, N-1]}, \quad (17g)$$

$$I_{N_q n_u} u_{i|k} \leq \bar{u}_\Sigma S_\tau^\top \delta_{i-\tau|k}, \quad \begin{cases} \forall \tau \in \mathbb{N}_{[1, \bar{\tau}]}, \\ \forall i \in \mathbb{N}_{[0, N-1]}. \end{cases} \quad (17h)$$

The cost function (17a) over a horizon $N = 8$, where $\mathbf{u}_k = (u_{0|k}, \dots, u_{N-1|k})$ and $\delta_k = (\delta_{0|k}, \dots, \delta_{N-1|k})$ denote the predicted input and activation signal sequences starting at time k , contains a tracking term to steer the ROI states z_k towards z_r , and a cost on $\bar{\epsilon}_{i|k}$ to incorporate the desired upper temperature bound as a soft constraint. The weights are chosen as

$$Q = \frac{1}{n_z} I_{n_z}, \quad f_\epsilon = 10, \quad (17i)$$

which are normalized with respect to the number of weighted variables for more intuitive tuning. The relative weighting by f_ϵ exceeds that of Q , since overheating is detrimental to treatment quality and hence its prevention is prioritized. The equality constraints (17b) capture the system dynamics in (15), with the initial condition given by (17c) determined by the observer (16). The inequality (17d) describes the upper bound in the form of a soft constraint, for which the slack variable $\bar{\epsilon}_{i|k}$ is restricted to be nonnegative by (17e). The power constraints on the individual sonication points are given by (17f), which in addition to the nonnegativity constraint (6) describe an upper power limit of $\bar{u} = 15$ W per sonication point for safety. Finally, the inequalities (17g)-(17h) incorporate the mode and setup time constraints (9), where in (17h) $\delta_{i-\tau|k} = \delta_{k+i-\tau}$ for $i < \tau$, as determined by the system's actuator path prior to the optimization at time k .




Fig. 5. The mean (solid) and maximum/minimum (dashed) temperature inside the ROI \mathcal{R} of the plant (black) and as estimated by the observer (gray).

Let $(\tilde{\delta}_k^*, \mathbf{u}_k^*)$ denote the pair of optimal activation signal and control input sequences found by solving (17) at time k given \hat{x}_k . As previously discussed, the mode and setup time constraints (17g)-(17h) ensure that the optimized input sequence \mathbf{u}_k^* corresponds to a physically realizable transducer trajectory, but that $\tilde{\delta}_k^*$ may not. Therefore, we feed the solution $(\tilde{\delta}_k^*, \mathbf{u}_k^*)$ to the admissibility assurance (14), from which we obtain an admissible sequence δ_k^* that also satisfies (17g)-(17h). Now, note that the pair $(\delta_k^*, \mathbf{u}_k^*)$ also minimizes (17), since \mathbf{u}_k^* is unaltered and (17a) contains no explicit cost on δ_k , and hence is an optimal admissible solution for the MI-MPC setup.

5. PERFORMANCE ANALYSIS

A simulation study has been performed to validate the temperature control performance of the MI-MPC for MR-HIFU hyperthermia. We initialize the plant, observer, and controller states at zero, corresponding to the monitored tissue to be at the blood temperature $T_b = 37$ °C before treatment. The temperature evolution of the voxels within the ROI of the plant (gray) and the observer/controller (black) are visualized in Fig. 5 using their mean (solid) and maximum/minimum values (dashed). It can be seen that after approximately 105 s the entire ROI is heated to 41 °C, which is the temperature above which the beneficial hyperthermia-related effects are adequately triggered. Furthermore, no overheating is observed, and the average ROI temperature converges to the optimum $T_r = 42$ °C. Using Matlab R2017b and Gurobi 8.1.1 on a laptop with Intel Core i7 @ 2.60 GHz CPU and 8 GB RAM, the computation times varied between 2.5 s and 5.5 s in steady state, with worst-case values around 11 s during heat-up. This indicates that real-time feasibility can likely be achieved, but may require dedicated hard- and software, or reducing the complexity of the MPC (e.g., by choosing a shorter horizon, or reducing the state and/or input dimensions).

In Fig. 6, the total acoustic input power per cell is shown, together with the desired-mode indicator σ_k . First, note that σ_k , which is directly related to δ_k , corresponds to a physically realizable transducer path, verifying that the proposed MI-MPC setup generates only admissible activation signal sequences. Regarding the input power, the controller initially requests maximum power for fast heat-up. Moreover, in this period it often heats a certain cell for several consecutive samples, thereby reducing actuator downtime, also contributing to achieving a short heat-




Fig. 6. The total acoustic power per cell, and the desired-mode indicator σ_k (red line).

up phase. For $t_k > 110$ s, however, each cell is only heated for one sample before continuing to the next cell, as this allows for maintaining an as homogeneous as possible temperature distribution throughout the ROI.

6. CONCLUSION

In this paper, an MI-MPC setup for large-area MR-HIFU hyperthermia is presented, which simultaneously optimizes the desired heating profile to be generated by the transducer, as well as the range-extending mechanical displacement of the transducer itself within a predefined discrete set of positions. For the latter, the proposed control design is able to explicitly take into account the point-to-point travel time of the transducer. Our modeling approach allows for including the transducer switching and travel times using intuitive specification, based on which the mixed-integer component of the controller follows systematically in a manner that directly yields a compact MIP. In a simulation study, the proposed controller is shown to adequately heat a large target area, thereby enabling the treatment of larger tumors. Valuable directions for future research include optimizing the treatment cell locations for control performance, improving the controller's robustness against model uncertainty, and reducing computation time.

ACKNOWLEDGEMENTS

The authors thank K.S. Mohan and J. van Wordragen for their contributions to this work, especially in building the software for the simulation study.

REFERENCES

Bemporad, A. and Morari, M. (1999). Control of systems integrating logic, dynamics, and constraints. *Automatica*, 35(3), 407–427.

Bing, C., Nofiele, J., et al. (2015). Localised hyperthermia in rodent models using an MRI-compatible high-intensity focused ultrasound system. *Int. J. Hyperth.*, 31(8), 813–822.

Datta, N.R., Ordóñez, S.G., et al. (2015). Local hyperthermia combined with radiotherapy and/or chemotherapy: recent advances and promises for the future. *Cancer Treat. Rev.*, 41(9), 742–53.

Datta, N.R., Puric, E., et al. (2016). Hyperthermia and Radiation Therapy in Locoregional Recurrent Breast Cancers: A Systematic Review and Meta-analysis. *Int. J. Radiat. Oncol.*, 94(5), 1073–1087.

Deenen, D.A., Maljaars, E., et al. (2018). Offset-free model predictive control for enhancing MR-HIFU hyperthermia in cancer treatment. *IFAC-PapersOnLine*, 51(20), 191–196.

Hijnen, N., Langereis, S., and Grull, H. (2014). Magnetic resonance guided high-intensity focused ultrasound for image-guided temperature-induced drug delivery. *Adv. Drug Deliv. Rev.*, 72, 65–81.

Hynynen, K. and Jones, R.M. (2016). Image-guided ultrasound phased arrays are a disruptive technology for non-invasive therapy. *Phys. Med. Biol.*, 61(17), R206–R248.

Ishihara, Y., Calderon, A., et al. (1995). A precise and fast temperature mapping using water proton chemical shift. *Magn. Reson. Med.*, 34(6), 814–823.

Issels, R.D., Lindner, L.H., et al. (2018). Effect of Neoadjuvant Chemotherapy Plus Regional Hyperthermia on Long-term Outcomes Among Patients With Localized High-Risk Soft Tissue Sarcoma. *JAMA Oncol.*, 4(4), 483–492.

Mallory, M., Gogineni, E., et al. (2016). Therapeutic hyperthermia: The old, the new, and the upcoming. *Crit. Rev. Oncol. Hematol.*, 97, 56–64.

Maloney, E. and Hwang, J.H. (2015). Emerging HIFU applications in cancer therapy. *Int. J. Hyperth.*, 31(3), 302–309.

Oei, A.L., Vriend, L.E., et al. (2015). Effects of hyperthermia on DNA repair pathways: One treatment to inhibit them all. *Radiat. Oncol.*, 10(1), 1–13.

Partanen, A., Yarmolenko, P.S., et al. (2012). Mild hyperthermia with magnetic resonance-guided high-intensity focused ultrasound for applications in drug delivery. *Int. J. Hyperth.*, 28(4), 320–336.

Pennes, H.H. (1948). Analysis of Tissue and Arterial Blood Temperatures in the Resting Human Forearm. *J. Appl. Physiol.*, 1(2), 93–122.

Philippe, M., Essick, R., et al. (2016). Stability of discrete-time switching systems with constrained switching sequences. *Automatica*, 72, 242–250.

Quesson, B., De Zwart, J.A., and Moonen, C.T.W. (2000). Magnetic resonance temperature imaging for guidance of thermotherapy. *J. Magn. Reson. Imaging*, 12(4), 525–533.

Rieke, V. and Butts Pauly, K. (2008). MR thermometry. *J. Magn. Reson. Imaging*, 27(2), 376–90.

Sebeke, L., Deenen, D.A., et al. (2019). Model predictive control for MR-HIFU-mediated, uniform hyperthermia. *Int. J. Hyperth.*, 36(1), 1040–1050.

Staruch, R.M., Hynynen, K., and Chopra, R. (2015). Hyperthermia-mediated doxorubicin release from thermosensitive liposomes using MR-HIFU: Therapeutic effect in rabbit Vx2 tumours. *Int. J. Hyperth.*, 31(2), 118–133.

Subramanian, K., Maravelias, C.T., and Rawlings, J.B. (2012). A state-space model for chemical production scheduling. *Comput. Chem. Eng.*, 47, 97–110.

Tillander, M., Hokland, S., et al. (2016). High intensity focused ultrasound induced in vivo large volume hyperthermia under 3D MRI temperature control. *Med. Phys.*, 43(3), 1539–1549.

Winter, L., Oberacker, E., et al. (2016). Magnetic resonance thermometry: Methodology, pitfalls and practical solutions. *Int. J. Hyperth.*, 32(1), 63–75.