

Offset-free model predictive control for enhancing MR-HIFU hyperthermia in cancer treatment [★]

D.A. Deenen ^{*} E. Maljaars ^{*} L. Sebeke ^{**} Bram 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 Cologne, Cologne, Germany

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

Abstract: In this paper, an offset-free MPC scheme is developed to improve performance and robustness of temperature control in magnetic resonance-guided high-intensity focused ultrasound hyperthermia cancer treatments. The proposed design effectively handles constraints and exploits the body's thermal behavior to achieve faster and more evenly distributed heating of the tumor leading to improved treatment quality with respect to the methods currently applied in clinics. The proposed scheme incorporates disturbance estimation to remove the steady-state offset resulting from plant-model mismatch, which for the given healthcare application is inevitable in practice due to the large physiological diversity in patients and tumors. Simulation results are presented to demonstrate this algorithm's capability to improve robust performance.

© 2018, IFAC (International Federation of Automatic Control) Hosting by Elsevier Ltd. All rights reserved.

Keywords: Healthcare, hyperthermia, offset-free MPC, cancer treatment, high-intensity focused ultrasound.

1. INTRODUCTION

Mild local hyperthermia involves the heating of a specific target volume inside the human body to temperatures of 39–45 °C for approximately 90 minutes. Numerous clinical trials have shown that hyperthermia is a potent adjuvant therapy for cancer treatment, see for example Datta et al. (2015); Mallory et al. (2016); Issels et al. (2018). It significantly increases the efficacy of primary methods such as chemo- and radiotherapy by sensitizing the treated tumorous tissue. Also, it enables the use of temperature-sensitive liposomes for heat-mediated drug release, allowing for more efficient delivery of the anticancer drugs specifically to the tumor region at lower systemic drug concentrations, as in Hijnen et al. (2014); Staruch et al. (2015). Most importantly, the results show that local hyperthermia does not increase toxicity in healthy tissue and the thereby induced side effects, making it highly appealing for clinical application to increase treatment success rates and improve quality of life for the patients in therapy.

A particularly interesting technology for hyperthermia is magnetic-resonance-guided high-intensity focused ultrasound (MR-HIFU), which enables heat delivery to internal tissues with spatial accuracy in the millimeter range, combined with real-time monitoring of the thermal response

via volumetric temperature maps obtained using an MRI scanner, see Maloney and Hwang (2015). This allows for a completely noninvasive procedure, thereby significantly contributing to patient well-being. However, realizing the desired temperature distribution and accurately maintaining it over the course of an entire treatment in a clinical setting is no simple task. Combined with the fact that the beneficial effects of thermal therapies, and of mild hyperthermia in particular, are closely related to the tissue temperatures truly achieved during treatment, see Franckena et al. (2009), this strongly motivates the development of *feedback* controllers for MR-HIFU.

Currently, (experimental) clinical applications of local hyperthermia by MR-HIFU typically involve a (mostly) pre-determined sonication plan, possibly extended with simple online MR-driven feedback algorithms. Examples include binary strategies to scale the sonication power and/or length of the heating interval (Enhölm et al. (2010); Partanen et al. (2012)), basic proportional-integral-derivative (PID)-based methods (Mougenot et al. (2009)) or some hybrid form of PID and bang-bang control (Staruch et al. (2011); Lin et al. (1990); Bing et al. (2015)). A major drawback of such designs, however, is their inability to take the body's future thermal behavior and the restrictive actuator constraints into account when computing the input, which negatively affects treatment quality and duration.

In this respect, model predictive control (MPC) offers superior potential for hyperthermia. This has been recognized by some researchers, of which the results are published in, for example, Arora et al. (2007) for the control of

[★] 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.

thermal dose in a single point based on a one-dimensional temperature distribution, De Bever et al. (2014) for thermal dose control in two-dimensional systems using heavily simplified models and a fixed sonication trajectory, and more recently in Sebeke et al. (2017) for temperature control in two-dimensional systems with online adaption of the heating location and power.

Although the initial results of online optimization-based control in the aforementioned works look promising, it is also clear that the full potential of MPC is far from reached. More sophisticated MPC setups specifically tailored to hyperthermia can fulfill the envisioned potential, resulting in more reliable, reproducible and high-performing 3D temperature control, which in turn may serve as an enabler for widespread and successful clinical use of MR-HIFU for hyperthermia. It is the objective of this paper to provide a significant step forward in the development of this societally highly relevant technology.

To this end, an offset-free MPC design is developed for planar heating using an MR-HIFU system. A Luenberger-type observer is employed to improve the MR thermometry measurements and to construct a disturbance estimate. The latter enables the controller to compensate for the effects of constant (or slowly varying) model errors, allowing for enhanced control performance in the clinic, where such model discrepancies are inevitable. Simulations are performed to demonstrate the developed model-based algorithm's potential in terms of achieving desired steady-state behavior despite severe plant-model mismatch.

The remainder of this paper is organized as follows. In Section 2 relevant aspects related to the clinical application of mild hyperthermia are discussed in more detail to clarify the challenges involved with feedback control for this application, and to motivate the desire to develop MPC for MR-HIFU. Next, Section 3 explains how the thermal model is obtained and Section 4 covers the controller design procedure. Simulations are performed to evaluate the algorithm's ability of removing the steady-state offset otherwise induced by model mismatch, the results of which are analyzed in Section 5. Finally, Section 6 summarizes the key achievements and corresponding observations.

2. PROBLEM DESCRIPTION

This section introduces some practical aspects concerning the hyperthermia treatment and setup, and discusses the motivation for using MPC in such an application.

2.1 MR-HIFU hyperthermia setup

Although the MPC setup we propose is generic in nature, in this paper it is designed for the Philips Sonalleve system shown in Fig. 1, which is a commercial MR-HIFU system already being used in clinics to treat uterine fibroids and for bone pain palliation therapy. It is a combined setup of a standard MRI scanner used for noninvasive thermography and a dedicated trolley-tabletop in which an MR-compatible HIFU transducer is integrated.

2.2 HIFU applicator

This system uses a phased-array HIFU transducer to generate the ultrasound waves. It consists of 256 elements



Fig. 1. Philips Sonalleve MR-HIFU therapy platform.




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

of which the phases and amplitudes can individually be chosen such that by interference a focal spot is created. By modulating these settings, which is referred to as electronic beam steering, the lateral location of the focal spot can be chosen anywhere on the focal plane within an 18 mm diameter circle around the transducer axis, as depicted in Fig. 2. Since the size of the focal spot is typically smaller than the tumor, it must be scanned through the treatment volume using beam steering or by moving the transducer. In this paper, we focus on beam steering in the focal plane using a predetermined set of discrete possible focus locations.

2.3 Motivation for offset-free MPC in hyperthermia

To fully benefit from the hyperthermia-induced effects such as increased blood flow and oxygenation, it is crucial to maintain temperatures above 41 °C in the region of interest (ROI), typically the tumor and possibly some adjacent tissue, during the entire 90 minute treatment. On the other hand, temperatures above 43 °C are undesired since they may cause reverse effects such as vascular shutdown (Hijnen et al. (2014)). Furthermore, temperature elevation outside the ROI must be avoided to prevent damage to healthy tissue.

MPC is considered superior to PID-based strategies for application in MR-HIFU hyperthermia, as it is able to reduce treatment time by exploiting beneficial (future) behavior, e.g., heat diffusion via conduction, and can explicitly incorporate actuator constraints such as the inability to actively remove heat from inside the body using HIFU. Unfortunately, model-based strategies are inherently accompanied by the possibility of modeling errors, which in the thermal models for hyperthermia treatment are highly likely to occur due to the high variability of spatially distributed and time/temperature-varying tissue parameters. Inadequately accounting for such mismatches may result in insufficient heating of the target region or in the overheating of healthy tissue, thereby significantly deteriorating treatment quality.




Fig. 3. Block scheme of envisioned fully automated MR-HIFU feedback loop for patient treatment.

Given this situation, we propose an MPC setup novel for hyperthermia inspired by an offset-free MPC algorithm as in Pannocchia and Rawlings (2003). Fig. 3 depicts the resulting feedback scheme, where the observer constructs the model-based temperature and disturbance estimates from the noninvasively obtained MR thermometry data, such that the MPC scheme can compute the optimal power distribution to be generated by the HIFU actuator.

3. THERMAL PROCESS MODELING

In this section, a model of the body's thermal response will be discussed, from which the state-space models used by the observer and MPC algorithm will be derived.

3.1 Bioheat model

The tissue's temperature change in response to the applied heating is modeled using Pennes' bioheat equation, see Pennes (1948), given by

$$\rho c \frac{\partial T(r, t)}{\partial t} = k \nabla^2 T(r, t) + Q(r, t) - w_b(r) c_b (T(r, t) - T_b), \quad (1)$$

where $T : \Omega \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is the temperature profile, such that $T(r, t)$ denotes the temperature at time $t \in \mathbb{R}_{\geq 0}$ and location $r = [r_x, r_y]^\top \in \Omega \subset \mathbb{R}^2$ with Ω being the patient domain in the focal plane. The parameters $\rho, c, k \in \mathbb{R}_{> 0}$ denote the tissue's volumetric mass density, specific heat capacity, and thermal conductivity, and $w_b : \Omega \rightarrow \mathbb{R}_{\geq 0}$ and $c_b \in \mathbb{R}_{> 0}$ the blood's perfusion coefficient and specific heat capacity, respectively. In this paper, the tissue is assumed to be homogeneous with respect to most parameters ($\rho, k, c,$ and c_b), but not the blood perfusion rate $w_b(r)$, which is assumed significant only in tissue in the direct vicinity of blood vessels, and zero elsewhere. Note, however, that our MPC setup can be extended to incorporate the fully inhomogeneous case as well.

The power deposition density described by $Q : \Omega \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ depends on the acoustic deposition intensity profile $F : \Omega \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ and linearly scales with the sonication power $P : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ according to

$$Q(r, t) = F(r, t)P(t), \quad r \in \Omega, t \in \mathbb{R}_{\geq 0}. \quad (2)$$

In reality, for a given focus location $r_f(t)$ at some time t , the intensity $F(r, t)$ also depends on factors such as the

acoustic properties of the intermediate tissue. In this paper, however, F is modeled by a radially symmetric two-dimensional Gaussian distribution centered around the focus location r_f with standard deviation $\sigma_f = 2.4$ mm

$$F(r, t) = \frac{\alpha}{2\pi\sigma_f^2} \exp\left(-\frac{\|r - r_f(t)\|^2}{2\sigma_f^2}\right), \quad (3)$$

with $r \in \Omega$, $t \in \mathbb{R}_{\geq 0}$, and scaling factor $\alpha \in \mathbb{R}_{> 0}$.

3.2 Discrete-time state-space model

Assuming zero out-of-plane interaction, (1)-(3) are discretized spatially using the central difference scheme on a 39×39 grid with a voxel size of 2×2 mm² and temporally using forward Euler with a sampling interval of $T_s = 2.5$ s, resulting in the discrete-time state-space model

$$x_{k+1} = A_p x_k + B_p u_k, \quad (4a)$$

$$y_k = x_k + v_k, \quad (4b)$$

with $k \in \mathbb{N}$ connecting to real time $t_k = kT_s$ and the states $x_k \in \mathbb{R}^n$, $n = 39^2 = 1521$ representing the voxel temperatures. We select 48 voxels within the ROI, at the center of which we allow sonication by steering the focal point r_f to this location. These locations are referred to as sonication points. The input $u_k \in \mathbb{R}^m$ with $m = 48$ is the applied sonication power averaged over the sampling interval at each of these sonication points. Each individual column of $B_p \in \mathbb{R}^{n \times m}$ captures the spatially discretized power deposition as in (2) and (3) corresponding to r_f at a single sonication point, and its temporally discretized relation to voxel temperature change according to (1).

The spatial discretization is chosen in such a manner that the centers of the voxels coincide with the points measured by MR thermometry. Also, the sampling interval T_s is larger than the time required for measurement, implying that at each discrete time $k \in \mathbb{N}$ a full new temperature map is available and thus the output $y_k \in \mathbb{R}^n$ can indeed be modeled as a full (noisy) state measurement as in (4b). In practice, the MR data is corrupted by noise, which for the purpose of this paper can be approximated by uncorrelated zero-mean Gaussian white noise $v_k \in \mathbb{R}^n$, $k \in \mathbb{N}$, with standard deviation $\sigma_v = \sqrt{0.2}$ °C.

4. CONTROLLER DESIGN

In this section, the state-space models used by the controller and observer are given first. Then, the control objectives as described in Subsection 2.3 are formulated as a constrained optimization problem.

4.1 Prediction model

To enable the controller to compensate for the steady-state offset induced by model mismatch or slowly varying disturbances, integral action is incorporated using disturbance estimation as described in Pannocchia and Rawlings (2003). The state-space representation of the prediction model is given by

$$x_{i+1|k} = A x_{i|k} + B u_{i|k} + d_{i|k}, \quad (5a)$$

$$d_{i+1|k} = d_{i|k}, \quad (5b)$$

$$y_{i|k} = x_{i|k}, \quad (5c)$$

where $x_{i|k}, u_{i|k}, y_{i|k} \in \mathbb{R}^n$ denote the predicted states, inputs, and outputs, respectively, at $i \in \mathbb{N}$ time steps ahead

of the prediction sequence's starting time k . Furthermore, $d_{i|k} \in \mathbb{R}^n$ denotes the virtual state disturbance, which is assumed to be constant over the prediction horizon. Note that we allow model mismatch, i.e., $A \neq A_p$ and $B \neq B_p$, and that we did not incorporate measurement noise in the prediction model since $\mathbb{E}(v_k) = 0$.

As proposed in Theorem 1 of Pannocchia and Rawlings (2003), we first verify that the nominal model (i.e., the model without additional disturbance $d_{i|k}$) is detectable (in fact, observable). Second, we evaluate the rank condition stated in the theorem, which using our disturbance model becomes

$$\text{rank} \begin{bmatrix} I - A & -I \\ I & 0 \end{bmatrix} = 2n, \quad (6)$$

to conclude that (5) is detectable, and therefore an asymptotically stable estimator exists given the proposed disturbance model. This property also guarantees that for a stable closed-loop system, a feasible steady-state setpoint, and constant disturbances, there will be zero offset between the realized and desired temperature distributions in case no constraints are active at steady state.

4.2 State and disturbance estimation

To facilitate this offset-free implementation, and to simultaneously obtain improved temperature estimates with respect to the noise-corrupted MR thermometry readings, we design the state and disturbance estimator

$$\hat{x}_k = A\hat{x}_{k-1} + Bu_{k-1} + \hat{d}_{k-1} + L_x(y_k - \hat{y}_k^-), \quad (7a)$$

$$\hat{d}_k = \hat{d}_{k-1} + L_d(y_k - \hat{y}_k^-), \quad (7b)$$

where

$$\hat{y}_k^- = A\hat{x}_{k-1} + Bu_{k-1} + \hat{d}_{k-1} \quad (8)$$

denotes the model-based measurement estimate at time k before applying the correction by output injection to obtain \hat{x}_k . Note that we choose the observer model (7) equal to the prediction model (5). Since we have full (noisy) state measurement, the system (5) is observable, allowing for arbitrary pole placement in (7) via L_x and L_d . We choose the diagonal matrices

$$L_x = \frac{1}{3}I_n, \quad L_d = \frac{1}{30}I_n, \quad (9)$$

which cause the state and disturbance estimation errors to converge to within 5% of their initial value in approximately 22.5 and 225 seconds, respectively.

4.3 Optimization problem

The temperature objectives have been discussed in Subsection 2.3 and are schematically depicted in Fig. 4 in cross-section perspective. $\underline{T} : \Omega \rightarrow \mathbb{R}$ and $\overline{T} : \Omega \rightarrow \mathbb{R}$ represent the (location-dependent) lower and upper temperature bounds, respectively, defining the desired temperature range (green) over the entire patient domain Ω . Under- and overheating occurs for temperatures within the (light) blue and (light) red areas in the figure, respectively. The concentric circular regions \mathcal{R} and \mathcal{S} denote the ROI and the region additionally including the tissue surrounding the ROI, i.e., $\mathcal{R} \subset \mathcal{S} \subset \Omega$. \underline{T} is nonzero only within \mathcal{R} where sufficient heating is desired. \overline{T} features an elevated plateau on \mathcal{S} , preventing reversal of the temperature-dependent mechanisms due to overheating, and has a

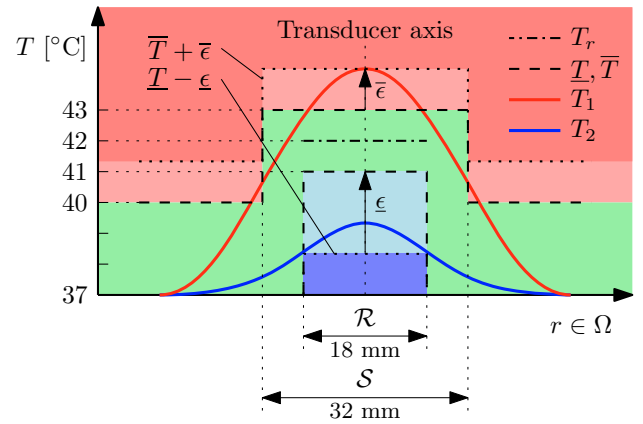


Fig. 4. Schematic cross-section of the temperature objectives for the circular regions \mathcal{R} and \mathcal{S} centered at the transducer axis. The maximum violations $\underline{\epsilon}$ and $\overline{\epsilon}$ are shown for some over- and underheated temperature distributions such that $T_1 \leq \overline{T} + \overline{\epsilon}$ and $T_2 \geq \underline{T} - \underline{\epsilon}$.

lower value on $\Omega \setminus \mathcal{S}$ to safeguard healthy tissue. Within the aforementioned temperature range, preference is given to a flat temperature distribution in the ROI, which is described by the reference temperature T_r .

To formulate these objectives mathematically, let us denote the elements of x corresponding to a point within \mathcal{R} by $x^{\mathcal{R}} = Hx \in \mathbb{R}^{n_{\mathcal{R}}}$, where $H \in \{0, 1\}^{n_{\mathcal{R}} \times n}$ is a matrix with exactly one 1 in each row. Furthermore, let $x^r \in \mathbb{R}^{n_{\mathcal{R}}}$ and $\underline{x}, \overline{x} \in \mathbb{R}^n$ denote the voxel-wise temperature reference and lower/upper bounds corresponding to the values of T_r, \underline{T} , and \overline{T} on the discrete voxel locations, respectively. The voxels' temperature deviations with respect to the reference in the ROI are then given by $x^{\mathcal{R}} - x^r$. The maximum violations of the temperature bounds follow from $\overline{\epsilon} = \|\max(x - \overline{x}, \mathbf{0})\|_{\infty} \geq 0$ and $\underline{\epsilon} = \|\max(\underline{x} - x, \mathbf{0})\|_{\infty} \geq 0$, where the maximum operator is used component-wise, and are collected in $\epsilon = [\underline{\epsilon}, \overline{\epsilon}]^T \in \mathbb{R}_{\geq 0}^2$.

The control objective is now formulated as the constrained optimization problem

$$\min_{u_{i|k}, \epsilon_{i|k}} \left(\sum_{i=0}^{N-1} \ell(x_{i|k}, u_{i|k}, \epsilon_{i|k}) + \ell(x_{N|k}, 0, \epsilon_{N|k}) \right) \quad (10a)$$

with, for some $i, k \in \mathbb{N}$, the stage cost

$$\ell(x_{i|k}, u_{i|k}, \epsilon_{i|k}) = (x_{i|k}^{\mathcal{R}} - x^r)^T Q (x_{i|k}^{\mathcal{R}} - x^r) + u_{i|k}^T R u_{i|k} + f_{\epsilon}^T \epsilon_{i|k}, \quad (10b)$$

subject to

$$\underline{x} - \mathbf{1}\underline{\epsilon}_{i|k} \leq x_{i|k} \leq \overline{x} + \mathbf{1}\overline{\epsilon}_{i|k}, \quad i \in [0, N], \quad (10c)$$

$$x_{i+1|k} = Ax_{i|k} + Bu_{i|k} + d_{i|k}, \quad i \in [0, N], \quad (10d)$$

$$d_{i+1|k} = d_{i|k}, \quad i \in [0, N-1], \quad (10e)$$

$$u_{i|k} \geq 0, \quad i \in [0, N-1], \quad (10f)$$

$$\mathbf{1}^T u_{i|k} \leq \overline{u}_{tot}, \quad i \in [0, N-1]. \quad (10g)$$

Inequality (10c), in which $\mathbf{1}$ denotes an all-ones vector of suitable dimension, captures the temperature bounds as soft constraints (note that $\epsilon_{i|k} = [\underline{\epsilon}_{i|k}, \overline{\epsilon}_{i|k}]^T$), (10d)-(10e) ensure satisfaction of the dynamics according to (5) with the prediction sequence's initial conditions given by $x_{0|k} = \hat{x}_k$ and $d_{0|k} = \hat{d}_k$, and finally (10f)-(10g) describe

the actuator constraints with $\overline{u_{tot}} = 60$ W being the maximum allowable power being applied to the treated region. Furthermore, a horizon of $N = 5$ is used, and the weights are chosen to be

$$Q = \frac{10}{n_{\mathcal{R}}} I, \quad R = \frac{0.01}{m} I, \quad f_{\epsilon} = \begin{bmatrix} 100 \\ 200 \end{bmatrix}, \quad (11)$$

which are normalized with respect to the number of corresponding variables for more intuitive balancing of the objectives' relative contribution to the cost function.

The critical importance of achieving temperatures within the desired range is expressed by the linear term (which is possible since $\epsilon \geq 0$ by definition) with high weighting. Using only the maximum violation, as opposed to penalizing each voxel individually, improves temperature homogeneity during the initial heat-up phase. To also encourage a flat temperature distribution (within the temperature bounds) in the ROI at steady state, a quadratic term with relatively lower cost is added to penalize this tracking error. In addition to the aforementioned control objectives, a small weight on the input power is incorporated to favor solutions that involve less heating.

5. PERFORMANCE ANALYSIS

A simulation study was carried out to verify the effectiveness of the proposed control algorithm in terms of compensating for the effects of plant-model mismatch for the MR-HIFU system under consideration. For reference, a simulation without model error is performed first, i.e., where $A = A_p$ and $B = B_p$, of which the results are shown in Fig. 5a. Initially, the entire patient volume is assumed to be 37 °C. The estimated temperature $\hat{x}^{\mathcal{R}}$ (gray) closely resembles its true counterpart $x^{\mathcal{R}}$ (black) in the ROI, as is shown in the figure by their lowest/highest (dashed) and average value (solid), enabling the controller to accurately heat the target region in a satisfactory manner. The current MPC's worst-case computation time is found to be 1 s, indicating it is sufficiently fast given $T_s = 2.5$ s.

Next, we perform a simulation in which using the nominal model results in temperature overestimation. To this end, the thermal conductivity k in (1) is intentionally underestimated and the power deposition intensity's scaling factor α in (3) is overestimated in the prediction and estimation models, both by a factor of 5, affecting A and B , respectively, in (5) and (7). Additionally, the observer and controller models are set to assume a homogeneous perfusion rate $w_b(r) = w_b$, thereby neglecting the local influence of blood vessels present in the plant. The results are shown in Fig. 5b. Initially, the disturbance estimator is disabled, i.e., $\hat{d}_k = 0$ and $L_d = 0$ for $k \in \mathbb{N}$ corresponding to $t_k < 300$ s. This causes a mismatch between $\hat{x}^{\mathcal{R}}$ and $x^{\mathcal{R}}$, resulting in a downward offset in the realized $x^{\mathcal{R}}$ with respect to its reference x^r . From $t_k = 300$ s (vertical dashed black line) onwards, the offset-free algorithm is active by setting L_d as in (9). At $t_k = 600$ s, after convergence of \hat{d} , the achieved temperature distribution coincides with that of Fig. 5a, illustrating how offset-free MPC can significantly improve hyperthermia treatment quality under model error.

This is more clearly illustrated in Fig. 6, where similarly to Fig. 4 cross-sections of the temperature objectives and distributions through the center of \mathcal{R} (at $r_x = 0.128$ m)




Fig. 5. The average (solid) and lowest/highest (dashed) voxel temperatures in the ROI of the plant $x^{\mathcal{R}}$ (black) and observer $\hat{x}^{\mathcal{R}}$ (gray) in case of (a) no plant-model mismatch, and (b) significant model error where disturbance estimation is enabled at $t = 300$ s.




Fig. 6. Voxel temperature x_k at cross-section through center of ROI without plant-model mismatch at $t_k = 600$ s (a), and with mismatch at $t_k = 300$ s (b) and $t_k = 600$ s (c), alongside the acoustic intensity around in-plane sonication points F_i . The gray areas indicate the cross-section voxels with increased perfusion.

are shown, alongside the (scaled) acoustic intensities $F_i = F(r_{f,i})$, $i = 1, \dots, 6$, from (3) with center locations $r_{f,i}$ at the six in-plane sonication points. The temperature profile corresponding to Fig. 5a at $t_k = 600$ s (black) shows the MPC's optimal performance without model error. It is within the desired temperature range, i.e., $\underline{x} \leq x \leq \overline{x}$, and tends to x^r in \mathcal{R} , however due to the system restrictions it is not a perfectly flat distribution. That is, the acoustic intensity profiles (blue) do not admit further heating of the areas with locally increased perfusion (gray) without causing additional overheating outside these areas. As previously found in Fig. 5b at $t_k = 300$ s, Fig. 6 shows that while introducing plant-model mismatch with disturbance estimation disabled




Fig. 7. Disturbance estimate \hat{d} at time $t_k = 600$ s, with additionally the edges of the circular \mathcal{R} (solid black) and \mathcal{S} (dashed black), sonication points (black dots), and voxels with increased perfusion (red squares).

causes insufficient heating (solid red), optimal performance is recovered by the offset-free algorithm after convergence at $t_k = 600$ s (dashed red).

The disturbance estimate after convergence at $t_k = 600$ s is visualized in Fig. 7. Indeed, overestimation occurs in the entire ROI due to the mismatch in k and α , which is compensated for by negative values in \hat{d} . Moreover, the figure shows that the disturbance estimator correctly identifies the voxels in which the locally increased perfusion is neglected (red outline), and thus is able to counteract the corresponding effects. In the figure, some voxels outside \mathcal{S} have a slightly positive disturbance estimate, which is a result of the propagation of measurement noise into \hat{d} via the output injection in (7).

6. CONCLUSION

In this paper, an offset-free MPC scheme is presented to enhance control performance of an MR-HIFU hyperthermia system used in cancer treatment. The proposed design achieves more accurate 2D temperature control than the techniques currently used in clinics. Specifically, the model-based algorithm is able to explicitly take into account (a) the restrictive input constraints and (b) the body's future thermal response, thereby enabling more optimal heating and shorter treatment times. Using disturbance estimation, the control design eliminates the steady-state offset otherwise resulting from modeling error, thus enhancing hyperthermia treatment quality. The MPC setup's effectiveness in terms of compensating for unmodeled perfusion and inaccurate conduction/heating parameters is demonstrated using simulation studies.

REFERENCES

Arora, D., Skliar, M., Cooley, D., and Roemer, R.B. (2007). Constrained Predictive Control of Thermal Therapies for Minimum-Time Delivery of Thermal Dose. *IEEE Trans. Control Syst. Technol.*, 15(6), 1030–1037.

Bing, C., Nofiele, J., Staruch, R., 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., Ordóñez, S.G., Gaip, U., et al. (2015). Local hyperthermia combined with radiotherapy and/or chemotherapy: Recent advances and promises for the future. *Cancer Treat. Rev.*, 41(9), 742–753.

De Bever, J., Todd, N., Payne, A., Christensen, D.A., and Roemer, R.B. (2014). Adaptive model-predictive controller for magnetic resonance guided focused ultrasound therapy. *Int. J. Hyperth.*, 30(7), 456–470.

Enholm, J., Kohler, M., Quesson, B., et al. (2010). Improved Volumetric MR-HIFU Ablation by Robust Binary Feedback Control. *IEEE Trans. Biomed. Eng.*, 57(1), 103–113.

Franckena, M., Fatehi, D., De Bruijne, M., et al. (2009). Hyperthermia dose-effect relationship in 420 patients with cervical cancer treated with combined radiotherapy and hyperthermia. *Eur. J. Cancer*, 45(11), 1969–1978.

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.

Issels, R.D., Lindner, L.H., Verweij, J., 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.

Lin, W.L., Roemer, R.B., and Hynynen, K. (1990). Theoretical and experimental evaluation of a temperature controller for scanned focused ultrasound hyperthermia. *Med. Phys.*, 17(4), 615–625.

Mallory, M., Gogineni, E., Jones, G.C., Greer, L., and Simone, C.B. (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.

Mougenot, C., Quesson, B., De Senneville, B.D., et al. (2009). Three-dimensional spatial and temporal temperature control with MR thermometry-guided focused ultrasound (MRgHIFU). *Magn. Reson. Med.*, 61(3), 603–614.

Pannocchia, G. and Rawlings, J.B. (2003). Disturbance models for offset-free model-predictive control. *AIChE J.*, 49(2), 426–437.

Partanen, A., Yarmolenko, P.S., Viitala, A., 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. (1948). Analysis of Tissue and Arterial Blood Temperatures in the Resting Human Forearm. *J. Appl. Physiol.*, 1(2), 93–122.

Sebeke, L., Luo, X., Heijman, E., et al. (2017). Prediction-based controller for MR-HIFU mediated hyperthermia. In *Eur. Symp. Focus. Ultrasound Ther.*, 23.

Staruch, R., Chopra, R., and Hynynen, K. (2011). Localised drug release using MRI-controlled focused ultrasound hyperthermia. *Int. J. Hyperth.*, 27(2), 156–171.

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.